MICROSOFT®

286 XENIX™ Operating System

Programmer's Guide

Programming Commands (CP)

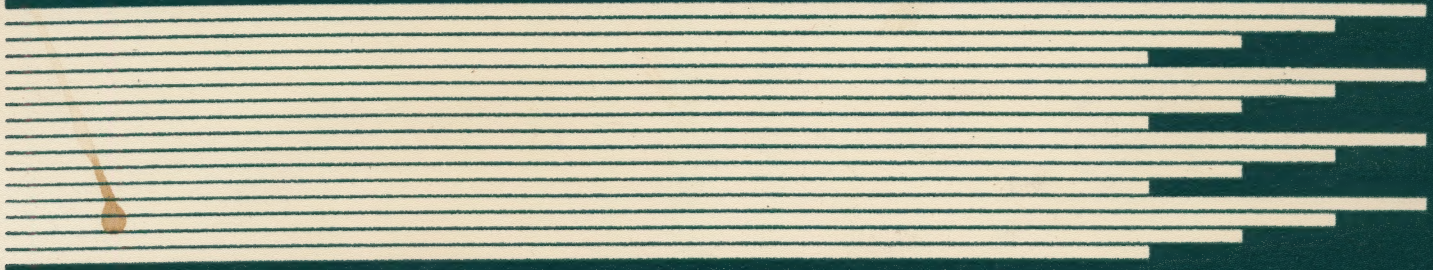

MICROSOFT®

MICROSOFT® XENIX™

Operating System

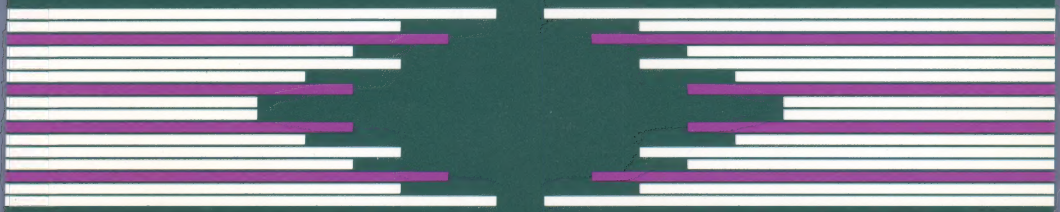

The 286 XENIX™ Operating System

Programmer's Guide

Information in this document is subject to change without notice and does not represent a commitment on the part of Microsoft Corporation. The software described in this document is furnished under a license agreement or nondisclosure agreement. The software may be used or copied only in accordance with the terms of the agreement. It is against the law to copy this software on magnetic tape, disk, or any other medium for any purpose other than the purchaser's personal use.

© Copyright Microsoft Corporation, 1984

Microsoft and the Microsoft logo are registered trademarks of Microsoft Corporation.

XENIX is a trademark of Microsoft Corporation.

Document Number: 8607E-300-00

Part Number: 091-092-011

Contents

1 Introduction

1.1 Overview 1-1
1.2 Creating C Language Programs 1-1
1.3 Creating Other Programs 1-2
1.4 Creating and Maintaining Libraries 1-2
1.5 Maintaining Program Source Files 1-2
1.6 Creating Programs With Shell Commands 1-3
1.7 Using This Guide 1-3
1.8 Notational Conventions 1-5

2 Cc: A C Compiler

2.1 Introduction 2-1
2.2 Invoking the C Compiler 2-1
2.3 Creating Programs From C Source Files 2-1
2.4 Creating Small, Middle, and Large Programs 2-5
2.5 Using Object Files and Libraries 2-7
2.6 Creating Smaller and Faster Programs 2-9
2.7 Preparing Programs for Debugging 2-11
2.8 Controlling the C Preprocessor 2-12
2.9 Error Messages 2-15
2.10 Using Advanced Options 2-17
2.11 Compiler Summary 2-21

3 Lint: A C Program Checker

3.1 Introduction 3-1
3.2 Invoking *lint* 3-1
3.3 Checking for Unused Variables and Functions 3-2
3.4 Checking Local Variables 3-3
3.5 Checking for Unreachable Statements 3-4
3.6 Checking for Infinite Loops 3-4
3.7 Checking Function Return Values 3-5
3.8 Checking for Unused Return Values 3-6
3.9 Checking Types 3-6
3.10 Checking Type Casts 3-7
3.11 Checking for Nonportable Character Use 3-7
3.12 Checking for Assignment of longs to ints 3-8
3.13 Checking for Strange Constructions 3-8
3.14 Checking for Use of Older C Syntax 3-9
3.15 Checking Pointer Alignment 3-10
3.16 Checking Expression Evaluation Order 3-10
3.17 Embedding Directives 3-11

3.18 Checking For Library Compatibility 3-12

4 Make: A Program Maintainer

4.1 Introduction 4-1
4.2 Creating a Makefile 4-1
4.3 Invoking Make 4-3
4.4 Using Pseudo-Target Names 4-4
4.5 Using Macros 4-5
4.6 Using Shell Environment Variables 4-8
4.7 Using the Built-In Rules 4-9
4.8 Changing the Built-in Rules 4-10
4.9 Using Libraries 4-12
4.10 Troubleshooting 4-13
4.11 Using Make: An Example 4-13

5 SCCS: A Source Code Control System

5.1 Introduction 5-1
5.2 Basic Information 5-1
5.3 Creating and Using S-files 5-5
5.4 Using Identification Keywords 5-13
5.5 Using S-file Flags 5-15
5.6 Modifying S-file Information 5-16
5.7 Printing from an S-file 5-20
5.8 Editing by Several Users 5-21
5.9 Protecting S-files 5-23
5.10 Repairing SCCS Files 5-25
5.11 Using Other Command Options 5-26

6 Adb: A Program Debugger

6.1 Introduction 6-1
6.2 Starting and Stopping Adb 6-1
6.3 Displaying Instructions and Data 6-4
6.4 Debugging Program Execution 6-13
6.5 Using the Adb Memory Maps 6-23
6.6 Miscellaneous Features 6-26
6.7 Patching Binary Files 6-32

7 As: An Assembler

7.1 Introduction 7-1
7.2 Command Usage 7-1
7.3 Lexical Conventions 7-2
7.4 Assembly Segments 7-3
7.5 Statements 7-4
7.6 Expressions 7-6

7.7 Assembler Directives 7-8
7.8 Machine Instructions 7-13
7.9 Addressing Modes 7-19
7.10 Diagnostics 7-23

8 Lex: A Lexical Analyzer

8.1 Introduction 8-1
8.2 Lex Source Format 8-2
8.3 Lex Regular Expressions 8-3
8.4 Invoking *lex* 8-4
8.5 Specifying Character Classes 8-5
8.6 Specifying an Arbitrary Character 8-6
8.7 Specifying Optional Expressions 8-6
8.8 Specifying Repeated Expressions 8-6
8.9 Specifying Alternation and Grouping 8-7
8.10 Specifying Context Sensitivity 8-7
8.11 Specifying Expression Repetition 8-8
8.12 Specifying Definitions 8-8
8.13 Specifying Actions 8-8
8.14 Handling Ambiguous Source Rules 8-12
8.15 Specifying Left Context Sensitivity 8-15
8.16 Specifying Source Definitions 8-17
8.17 Lex and Yacc 8-18
8.18 Specifying Character Sets 8-22
8.19 Source Format 8-23

9 Yacc: A Compiler-Compiler

9.1 Introduction 9-1
9.2 Specifications 9-4
9.3 Actions 9-6
9.4 Lexical Analysis 9-8
9.5 How the Parser Works 9-10
9.6 Ambiguity and Conflicts 9-14
9.7 Precedence 9-19
9.8 Error Handling 9-22
9.9 The Yacc Environment 9-24
9.10 Preparing Specifications 9-25
9.11 Input Style 9-25
9.12 Left Recursion 9-26
9.13 Lexical Tie-ins 9-27
9.14 Handling Reserved Words 9-27
9.15 Simulating Error and Accept in Actions 9-28
9.16 Accessing Values in Enclosing Rules 9-28
9.17 Supporting Arbitrary Value Types 9-29
9.18 A Small Desk Calculator 9-30
9.19 Yacc Input Syntax 9-32

9.20 An Advanced Example 9-34
9.21 Old Features 9-40

A The C-Shell

A.1 Introduction A-1
A.2 Invoking the C-shell A-1
A.3 Using Shell Variables A-2
A.4 Using the C-Shell History List A-4
A.5 Using Aliases A-7
A.6 Redirecting Input and Output A-8
A.7 Creating Background and Foreground Jobs A-9
A.8 Using Built-In Commands A-9
A.9 Creating Command Scripts A-11
A.10 Using the argv Variable A-11
A.11 Substituting Shell Variables A-12
A.12 Using Expressions A-14
A.13 Using the C-Shell: A Sample Script A-15
A.14 Using Other Control Structures A-18
A.15 Supplying Input to Commands A-19
A.16 Catching Interrupts A-20
A.17 Using Other Features A-20
A.18 Starting a Loop at a Terminal A-21
A.19 Using Braces with Arguments A-22
A.20 Substituting Commands A-22
A.21 Special Characters A-23

B C Language Portability

B.1 Introduction B-1
B.2 Program Portability B-2
B.3 Machine Hardware B-2
B.4 Compiler Differences B-7
B.5 Program Environment Differences B-12
B.6 Portability of Data B-12
B.7 Lint B-13
B.8 Byte Ordering Summary B-13

C Building a Communication System

C.1 Introduction C-1
C.2 What You Need C-1
C.3 Installing the Modem C-2
C.4 Creating a Dial-in Line C-5
C.5 Creating a Dial-out Line C-5
C.6 Installing a Uucp System C-8
C.7 Maintaining the System C-17
C.8 Details of Operation C-20

C.9 Creating a New dial Program C-30

D M4: A Macro Processor

D.1 Introduction D-1
D.2 Invoking *m4* D-1
D.3 Defining Macros D-2
D.4 Quoting D-3
D.5 Using Arguments D-5
D.6 Using Arithmetic Built-ins D-6
D.7 Manipulating Files D-7
D.8 Using System Commands D-7
D.9 Using Conditionals D-8
D.10 Manipulating Strings D-8
D.11 Printing D-10

Chapter 1
Introduction

1.1 Overview 1-1

1.2 Creating C Language Programs 1-1

1.3 Creating Other Programs 1-2

1.4 Creating and Maintaining Libraries 1-2

1.5 Maintaining Program Source Files 1-2

1.6 Creating Programs With Shell Commands 1-3

1.7 Using This Guide 1-3

1.8 Notational Conventions 1-5

1.1 Overview

This guide explains how to use the XENIX Software Development system to create and maintain C language and assembly language programs. The system provides a broad spectrum of programs and commands to help you design and develop applications and system software. These programs and commands enable you to create C and assembly language programs for execution on the XENIX system. They also let you debug these programs, automate their creation, and maintain different versions of the programs you develop.

The following sections introduce the programs and commands of the XENIX Software Development System, and explain the steps you can take to develop programs for the XENIX system. Most of the programs and commands in these introductory sections are fully explained later in this guide. Some commands mentioned here are part of the XENIX Timesharing System. These are explained in the XENIX *User's Guide* and XENIX *Operations Guide*.

1.2 Creating C Language Programs

All C language programs start as a collection of C program statements in a source file. The XENIX system provides a number of text editors that let you create source files easily and efficiently. The most convenient editor is the screen-oriented editor **vi**. **Vi** provides many editing commands that let you easily insert, replace, move, and search for text. All commands can be invoked from command keys or from a command line. **Vi** also has a variety of options that let you modify its operation.

Once a C language program has been written to a source file, you can create an executable program by using the **cc** command. The **cc** command invokes the XENIX C compiler which compiles the source file. This command also invokes other XENIX programs to prepare the compiled program for execution.

You can debug an executable C program with the XENIX debugger **adb**. **Adb** provides a direct interface to the machine instructions that make up an executable program.

If you wish to check a program before compiling it, you can use **lint**, the XENIX C program checker. **Lint** checks the content and construction of C language programs for syntactical and logical errors. It also enforces a strict set of guidelines for proper C programming style. **Lint** is normally used in the early stages of program development to check for illegal and improper usage of the C language.

Another way to check a program is with **cb**, the XENIX C program beautifier. **Cb** improves readability of C programs, making detection of logical errors easier.

1.3 Creating Other Programs

The C programming language can meet the needs of most programming projects. In cases where finer control of execution is required, you may create assembly language programs using the XENIX assembler **as**. **As** assembles source files and produces relocatable object files that can be linked to your C language programs with **ld**. The **ld** program is the XENIX linker. It links relocatable object files created by the C compiler or assembler to produce executable programs. Note that the **cc** command automatically invokes the linker and the assembler, so use of either **as** or **ld** is optional.

You can create source files for lexical analyzers and parsers using the program generators **lex** and **yacc**. Lexical analyzers are used in programs to pick patterns out of complex input and convert these patterns into meaningful values or tokens. Parsers are used in programs to convert meaningful sequences of tokens and values into actions. The **lex** program is the XENIX lexical analyzer generator. It generates lexical analyzers, written in C program statements, from given specification files. The **yacc** program is the XENIX parser generator. It generates parsers, written in C program statements, from given specification files. **Lex** and **yacc** are often used together to make complete programs.

You can preprocess C and assembly language source files, or even **lex** and **yacc** source files using the **m4** macro processor. The **m4** program performs several preprocessing functions, such as converting macros to their defined values and including the contents of files into a source file.

1.4 Creating and Maintaining Libraries

You can create libraries of useful C and assembly language functions and programs using the **ar** and **ranlib** programs. **Ar**, the XENIX archiver, can be used to create libraries of relocatable object files. **Ranlib**, the XENIX random library generator, converts archive libraries to random libraries and places a table of contents at the front of each library.

The **lorder** command finds the ordering relation in an object library. The **tsort** command topologically sorts object libraries so that dependencies are apparent.

1.5 Maintaining Program Source Files

You can automate the creation of executable programs from C and assembly language source files and maintain your source files using the **make** program and the SCCS commands.

The **make** program is the XENIX program maintainer. It automates the steps required to create executable programs, and provides a mechanism for

ensuring up-to-date programs. It is used with medium-scale programming projects.

The Source Code Control (SCCS) commands let you maintain different versions of a single program. The commands compress all versions of a source file into a single file containing a list of differences. These commands also restore compressed files to their original size and content.

Many XENIX commands let you carefully examine a program's source files. The **ctags** command creates a tags file so that C functions can be quickly found in a set of related C source files. The **mkstr** command creates an error message file by examining a C source file.

Other commands let you examine object and executable binary files. The **nm** command prints the list of symbol names in a program. The **hd** command performs a hexadecimal dump of given files, printing files in a variety of formats, one of which is hexadecimal. The **size** command reports the size of an object file. The **strings** command finds and prints readable text (strings) in an object or other binary file. The **strip** command removes symbols and relocation bits from executable files. The **sum** command computes a checksum value for a file and a count of its blocks. It is used in looking for bad spots in a file and for verifying transmission of data between systems. The **xstr** command extracts strings from C programs to implement shared strings.

1.6 Creating Programs With Shell Commands

In some cases, it is easier to write a program as a series of XENIX shell commands than it is to create a C language program. Shell commands provide much of the same control capability as the C language, and give direct access to all the commands and programs normally available to the XENIX user.

The **csh** command invokes the C-shell, a XENIX command interpreter. The C-shell interprets and executes commands taken from the keyboard or from a command file. It has a C-like syntax which makes programming in this command language easy. It also has an aliasing facility, and a command history mechanism.

1.7 Using This Guide

This guide is intended for programmers who are familiar with the C programming language and with the XENIX system.

Chapter 1 introduces the XENIX software development programs provided with this package.

Chapter 2 explains how to compile C language programs using the **cc** command.

Chapter 3 explains how to check C language programs for syntactical and semantical correctness using the C program checker **lint**.

Chapter 4 explains how to automate the development of a program or other project using the **make** program.

Chapter 5 explains how to control and maintain all versions of a project's source files using the SCCS commands.

Chapter 6 explains how to debug C and assembly language programs using the XENIX debugger **adb**.

Chapter 7 explains how to assemble assembly language programs using the XENIX assembler **as**.

Chapter 8 explains how to create lexical analyzers using the program generator **lex**.

Chapter 9 explains how to create parsers using the program generator **yacc**.

Appendix A explains how to use the C shell, a command interpreter that provides greater flexibility and more power than the standard XENIX shell, **sh**.

Appendix B explains how to write C language programs that can be compiled on other XENIX systems.

Appendix C explains how to build a communication link to other XENIX systems.

Appendix D explains how to use to create and process macros using the **m4** macro processor.

C language programmers should read Chapters 2, 3, and 6 for an explanation of how to compile and debug C language programs.

Assembly language programmers should read Chapter 7 for an explanation of the XENIX assembler and Chapter 6 for an explanation of how to debug programs.

Programmers who wish to automate the compilation process of their programs should read Chapter 4 for an explanation of the **make** program. Programmers who wish to organize and maintain multiple versions of their programs should read Chapter 5 for an explanation of the Source Code Control System (SCCS) commands.

Special project programmers who need a convenient way to produce lexical analyzers and parsers should read Chapters 8 and 9 for explanations of the **lex** and **yacc** program generators.

1.8 Notational Conventions

This guide uses a number of special symbols to describe the syntax of XENIX commands. The following is a list of these symbols and their meaning.

[ ]	Brackets indicate an optional command argument.
...	Ellipses indicate that the preceding argument may be repeated one or more times.
SMALL	Small capitals indicate a key to be pressed.
bold	Boldface characters indicate a command or program name.
italics	Italic characters indicate a placeholder for a command argument. When typing a command, a placeholder must be replaced with an appropriate filename, number, or option.

Chapter 2
Cc: A C Compiler

2.1 Introduction 2-1

2.2 Invoking the C Compiler 2-1

2.3 Creating Programs From C Source Files 2-1
- 2.3.1 Compiling a C Source File 2-2
- 2.3.2 Compiling Several Source Files 2-3
- 2.3.3 Naming the Output File 2-4

2.4 Creating Small, Middle, and Large Programs 2-5
- 2.4.1 Creating Small Model Programs 2-6
- 2.4.2 Creating Pure-Text Small Model Programs 2-6
- 2.4.3 Creating Middle Model Programs 2-6
- 2.4.4 Creating Large Model Programs 2-7

2.5 Using Object Files and Libraries 2-7
- 2.5.1 Creating Object Files 2-7
- 2.5.2 Creating Programs From Object Files 2-8
- 2.5.3 Linking a Program to Functions In Libraries 2-8

2.6 Creating Smaller and Faster Programs 2-9
- 2.6.1 Creating Optimized Object Files 2-9
- 2.6.2 Stripping the Symbol Table 2-10
- 2.6.3 Removing Stack Probes From a Program 2-11

2.7 Preparing Programs for Debugging 2-11
- 2.7.1 Producing an Assembly Language Listing 2-11
- 2.7.2 Profiling a Program 2-12

2.8 Controlling the C Preprocessor 2-12
- 2.8.1 Defining a Macro 2-13
- 2.8.2 Defining Include Directories 2-14
- 2.8.3 Ignoring the Default Include Directories 2-14
- 2.8.4 Saving a Preprocessed Source File 2-14

2.9 Error Messages 2-15
2.9.1 C Compiler Messages 2-15
2.9.2 Setting the Level of Warnings 2-16

2.10 Using Advanced Options 2-17
2.10.1 Creating Programs From Assembly Language Source Files 2-17
2.10.2 Using the near and far Keywords 2-17
2.10.3 Changing Word Order in Programs 2-18
2.10.4 Setting the Stack Size 2-19
2.10.5 Using Modules, Segments, and Groups 2-19

2.11 Compiler Summary 2-21
2.11.1 Cc Options 2-21
2.11.2 Memory Models 2-22
2.11.3 Pointer and Integer Sizes 2-23
2.11.4 Segment and Module Names 2-23

2.1 Introduction

This chapter explains how to use the cc command. In particular, it explains how to

- Compile C language source files
- Choose a memory model for a program
- Use object files and libraries with a program
- Create smaller and faster programs
- Prepare C programs for debugging
- Control the C preprocessor

It also describes the error and warning messages generated by the C compiler, and explains how to use advanced features of the cc command to make customized programs.

This chapter assumes that you are familiar with the C programming language, and that you can create C program source files using a XENIX text editor. For a description of the C language, see the XENIX *Microsoft C Reference Manual*.

2.2 Invoking the C Compiler

The cc command has the form

cc [*option*] ... *filename* ...

where *option* is a command option, and *filename* is the name of a C language source file, an assembly language source file, an object file, or an archive library. You may give more than one option or filename, if desired, but must separate each item with one or more spaces.

The cc command options let you control and modify the tasks performed by the command. For example, you can direct cc to perform optimization or create an assembly listing file. The options also let you specify additional information about the compilation, such as which program libraries to examine and what the name of the executable file should be. Many options are described in the following sections. For a complete description of all options, see *cc*(CP) in the XENIX *Reference Manual*.

2.3 Creating Programs From C Source Files

The cc command is normally used to create executable programs from C

language source files. A file's contents are identified by the filename extension. C source files must have the extension ".c".

The cc command can create executable programs only from source files that make up a complete C program. In XENIX, a complete program must have one (and only one) function named "main". This function becomes the entry point for program execution. The "main" function may call other functions as long as they are defined within the program or are part of the C standard library. The standard C library is described in the XENIX *Library Guide*.

2.3.1 Compiling a C Source File

You can compile a C source file by giving the name of the file when you invoke the cc command. The command compiles the statements in the file, then copies the executable program to the default output file *a.out*.

To compile a source program, type

```
cc filename
```

where *filename* is the name of the file containing the program. The program must be complete, that is, it must contain a "main" program function. It may also contain calls to functions explicitly defined by the program or by the standard C library.

For example, assume that the following program is stored in the file named *main.c*.

```
#include <stdio.h>

main ()
{
        int x,y;

        scanf("%d  %d", &x, &y);
        printf("%d\n", x+y);
}
```

To compile this program, type:

```
cc main.c
```

The command first invokes the C preprocessor, which adds the statements in the file */usr/include/stdio.h* to the beginning of the program. It then compiles these statements and the rest of the program statements. Next, the command links the program with the standard C library, which contains the object files for the *scanf* and *printf* functions. Finally, it copies the program to the file *a.out*.

You can execute the new program by typing

```
a.out
```

The program waits until you enter two numbers, then prints their sum. For example, if you type "3 5" the program displays "8".

2.3.2 Compiling Several Source Files

Large source programs are often split into several files to make them easier to understand, update and edit. You can compile such a program by giving the names of all the files belonging to the program when you invoke the cc command. The command reads and compiles each file in turn, then links all object files together, and copies the new program to the file *a.out*.

To compile several source files, type

```
cc filename ...
```

where each *filename* is separated from the next by at least one space. One of these files (and only one) must contain a "main" function. The others may contain functions that are called by this "main" function or by other functions in the program. The files must not contain calls to functions that are not explicitly defined by the program or by the standard C library.

For example, suppose the following main program function is stored in the file *main*.

```
#include <stdio.h>
extern int add();

main ()
{
        int x,y,z;

        scanf ("%d  %d", &x, &y);
        z = add (x, y);
        printf ("%d \n", z);
}
```

Assume that the following function is stored in the file *add.c*.

```
add (a, b)
int a, b;
{
        return (a + b);
}
```

You can compile these files and create an executable program by typing:

```
cc main.c add.c
```

The command compiles the statements in *main.c*, then compiles the statements in *add.c*. Finally, it links the two together (along with the standard C library) and copies the program to *a.out*. This program, like the program in the previous section, waits for two numbers, then prints their sum.

Since the **cc** command cannot keep track of more than one compiled file at a time, when several source files are compiled at a time, the command creates object files to hold the binary code generated for each source file. These object files are then linked to create an executable program. The object files have the same basename as the source files, but are given the ".o" file extension. For example, when you compile the two source files above, the compiler produces the object files *main.o* and *add.o*. These files are permanent files, i.e., the command does not delete them after completing its operation. Note that the command also creates an object file if only one source file is compiled.

2.3.3 Naming the Output File

You can give the executable program file any valid filename by using the –o (for "output") option. The option has the form

```
–o filename
```

where *filename* is a valid filename or pathname. If a filename is given, the program file is stored in the current directory. If a full pathname is given, the file is stored in the given directory. If that file already exists, its contents are replaced with the new executable program.

For example, the command

```
cc –o addem main.c add.o
```

causes the compiler to create an executable program file *addem* from the source file *main.c* and object file *add.o*. You can execute this program by typing:

```
addem
```

Note that the –o option does not affect the existing *a.out* file. This means that

the cc command does not change the current contents of *a.out* if the –o option has been given.

2.4 Creating Small, Middle, and Large Programs

The cc command lets you create programs of a variety of sizes and purposes using the –Ms, –Mm, –Ml, and –i options. These options define the size of a given program by defining the number of segments in physical memory to be allocated for your program's use. They also determine how the system loads the program for execution.

The cc command allows the creation of programs in four different memory models: impure-text small model, pure-text small model, middle model, and large model. Each model defines a different type of program structure and storage.

Impure-text small model programs are typically C programs that are short or have a limited purpose. These programs must not exceed 64 Kbytes.

Pure-text small model programs are typically short programs that are intended to be invoked by many users. Pure-text programs can occupy up to 128 Kbytes, but no more than 64 Kbytes each is permitted for either instructions or data. Unlike small model programs, the system loads only one copy of a pure-text program's instructions into memory, no matter how many times it has been invoked. As long as this copy stays in memory, the system simply loads a new copy of the data for each new invocation of the program. It then keeps each copy of data separate, while sharing the instructions among the different invocations. Pure-text programs save valuable memory space that would otherwise be wasted by small model programs.

Middle model programs are typically C programs, that have a large number of program statements but a relatively small amount of data. Program instructions can be any size, but program data must not exceed 64 Kbytes.

Large model programs are typically very large C programs which use a large amount of data storage during normal processing. Program instructions and data may have any size, except that the program must not contain arrays or structures that exceed 64 Kbytes.

C programs in memory consist of the actual machine instructions created from the program's source statements, and the several bytes of binary data storage created for the program's variables. The data storage also contains the stack used by the program for temporary storage during execution. The XENIX system stores the instructions and data in one or more segments of physical memory. Each segment is 64 Kbytes long. Thus, the maximum allowable size for any program depends on how many segments allocated for it when compiled.

The following sections describe how to use the –M and –i options to create programs with a specific number of segments. They also describe how to create pure-text programs for execution by multiple users.

2.4.1 Creating Small Model Programs

You can create a small model program by using the –Ms option. This option directs cc to create a program that occupies a single segment when loaded into physical memory. To create a small model program, type

cc –Ms *filename*

where *filename* is the name of the program you wish to compile.

The cc command creates small model programs by default when you do not otherwise specify a program model. Thus, the –Ms option is not required.

2.4.2 Creating Pure-Text Small Model Programs

You can create a pure-text small model program by combining the –i and –Ms options. The –i option directs cc to create separate memory segments for the instructions and data of a small model program. To create a pure-text program, type

cc –Ms –i *filename*

where *filename* is the name of the file source program to be compiled. Since cc creates small model programs by default, only the –i option is required.

2.4.3 Creating Middle Model Programs

You can create a middle model program by using the –Mm option. This option creates one segment for the data of the program, and one or more segments for the instructions. To create a middle model program, type

cc –Mm *filename* ...

where *filename* is the name of the source file to be compiled. When creating a program, the compiler attempts to fit as many instructions into a segment (up to 64 Kbytes) as possible. If the program is larger than 64 Kbytes, the –NGT option must be used to create new segments (see the section "Using Modules, Segments, and Groups" given later in this chapter).

Middle model programs are pure in the sense that the system never loads more than one copy of the program's instructions into memory at one time. This means the –i option, used with pure-text small model programs, is not required for middle model programs.

2.4.4 Creating Large Model Programs

You can create large model programs by using the –Ml option. This option directs cc to create multiple segments for both instructions and data. To create a large model program, type

cc –Ml *filename*

where *filename* is the name of a source file to be compiled.

As with middle model programs, the compiler attempts to fit as many instructions into a segment as possible. If a program's instructions or data is greater than 64 Kbytes each, the –NGT and –NGD options must be used to create new segments (see the section "Using Modules, Segments, and Groups" given later in this chapter).

Like middle model programs, large model programs are considered to be pure.

2.5 Using Object Files and Libraries

The cc command lets you save useful functions as object files, and use these object files to create programs at a later time. Object files contain the compiled or assembled instructions of your source file, so they save you the time and trouble of recompiling the functions each time you need them. All object files created by cc have the file extension ".o".

The cc command also lets you use functions found in XENIX system libraries, such as the standard C library or the screen processing library *curses*. To use these functions, you simply supply the name of the library containing them. In some cases, such as for the standard C library, cc accesses the library automatically and no explicit naming is required.

For convenience, you can create your own libraries with the **ar** and **ranlib** commands. These commands, described in section CP of the XENIX *Reference Manual*, copy your useful object files to a library file, and prepare the file for use by the cc command. You can access the library like any other library in the system if you copy it to the */lib* directory.

2.5.1 Creating Object Files

You can create an object file from a given source file by using the –c (for "compile") option. This option directs cc to compile the source file without creating a final program. The option has the form

–c *filename* ...

where *filename* is the name of the source file. You may give more than one

filename if you wish. Make sure each name is separated from the next by a space.

To make object files for the source files *add.c* and *mult.c*, type:

```
cc -c add.c mult.c
```

This command compiles each file and copies the compiled source files to the object files *add.o* and *mult.o*. It does not link these files; no executable program is created.

The -c option is typically used to save useful functions for programs to be developed later. Once a function is in an object file it may be used as is, or saved in a library file and accessed like other library functions, as described in the following sections.

Note that the cc command automatically creates object files for each source file in the command line. Unless the -c option is given, however, it will also attempt to link these files, even if they do not form a complete program.

2.5.2 Creating Programs From Object Files

You can use the cc command to create executable programs from one or more object files, or from a combination of object files and C source files. The command compiles the source files (if any), then links the compiled source files with the object files to create an executable program.

To create a program, give the names of the object and source files you wish to use. For example, if the source file *main.c* contains calls to the functions *add* and *mult* (saved in the object files *add.o* and *mult.o*), you can create a program by typing:

```
cc main.c add.o mult.o
```

In this case, *main.c* is compiled, then linked with *add.o* and *mult.o* to create the executable file *a.out*.

2.5.3 Linking a Program to Functions In Libraries

You can link a program to functions in a library by using the -l (for "library") option. The option directs cc to search the given library for the functions called in the source file. If the functions are found, the command links them to the program file.

The option has the form

```
cc -lname
```

where *name* is a shortened version of the library's actual filename (see *Intro*(S) in the XENIX *Reference Manual* for a list of names). Spaces between the name and option are optional. The linker searches the */lib* directory for the library. If not found, it searches the */usr/lib* directory.

For example, the command

```
cc main.c -lcurses
```

links the library */lib/libcurses.a* to the source file *main.c*.

A library is a convenient way to store a large collection of object files. The XENIX system provides several libraries, the most common of which is the standard C library. Functions in this library are automatically linked to your program whenever you invoke the compiler. Other libraries, such as *libcurses.a*, must be explicitly linked using the -l option. The XENIX libraries and their functions are described in detail in the XENIX *Library Guide*.

In general, the cc command does not search a library until the -l option is encountered, so the placement of the option is important. The option must follow the names of any source files containing calls to functions in the given library. In general, all library options should be placed at the end of the command line, after all source and object files.

2.6 Creating Smaller and Faster Programs

You can create smaller and faster C programs by using the optimizing options available with the cc command. These options reduce the size of a compiled program by removing unnecessary or redundant instructions or unnecessary symbol information. Smaller programs usually run faster and save valuable space.

2.6.1 Creating Optimized Object Files

You can create an optimized object file or an optimized program from a given source file by using the -O (for "optimize") option. This option reduces the size of the object file or program by removing unnecessary instructions. For example, the command

```
cc -O main.c
```

creates an optimized program from the source file *main.c*. The resulting object file or program is smaller (in bytes) than if the source had been compiled without the option. A smaller object file usually means faster execution.

The -O option applies to source files only; existing object files are ignored if included with this option. The option must appear before the names of the files you wish to optimize. For example, the command

```
cc -O add.c main.c
```

optimizes *main.c* and *add.c*.

You may combine the -O and -c options to compile and optimize source files without linking the resulting object files. For example, the command

```
cc -O -c main.c add.c
```

creates separate optimized object files from the source files *main.c* and *add.c*.

Although optimization is very useful for large programs, it takes more time than regular compilation. In general, it should be used in the last stage of program development, after the program has been debugged.

2.6.2 Stripping the Symbol Table

You can re luce the size of a program's executable file by using the -s and -x options. These options direct cc to remove items from the symbol table. The symbol table contains information about code relocation and program symbols and is used by the XENIX debugger *adb* to allow symbolic references to variables and functions when debugging. The information in this table is not required for normal execution, and should be removed when the program has been completely debugged.

The -s option strips the entire table, leaving machine instructions only. For example, the command

```
cc -s main.c add.c
```

creates an executable program that contains no symbol table. It also creates the object files *main.o* and *add.o* which contain no symbol tables.

The -x option strips all nonglobal symbols from the file including the names of local functions and variables, but excluding externally declared items. The command

```
cc -x main.o add.o
```

creates an executable program with global symbols, but only if the object files *main.o* and *add.o* have symbol tables.

The -s and -x options may be combined with the -O option to create an optimized and stripped program. Note that you can also strip a program with the XENIX command **strip**. See the XENIX *Reference Manual* for details.

2.6.3 Removing Stack Probes From a Program

You can reduce the size of a program slightly by using the –K option to remove all stack probes. A stack probe is a short routine called by a function to check the program stack for available space. The probes are not needed if the program makes very few function calls or has unlimited stack space.

To remove the stack probes from the program *main.c*, type

```
cc –K main.c
```

Although this option, when combined with the –O option, makes the smallest possible program, it should be used with great care. Removing stack probes from a program whose stack use is not well known can cause execution errors.

2.7 Preparing Programs for Debugging

The cc command provides a variety of options to prepare a program that is under development for debugging. These options range from creating an assembly language listing of the program, for use with the XENIX debugger **adb**, to adding routines for profiling the execution of a program.

2.7.1 Producing an Assembly Language Listing

You can direct the compiler to generate an assembly language listing of your compiled source file by using the –S and –L options. The –S option creates an assembly language listing that is suitable as input to the XENIX assembler **as**. The –L option creates a listing that shows assembled code, as well as instructions. The file created by –S is given the file extension ".s"; the file created by –L is given ".L".

Assembly language listing files are typically used by programmers who wish to debug their program with **adb**. Since **adb** recognizes machine instructions instead of the actual source statements in your program, a programmer needs an assembly language listing for accurate debugging.

To create an assembly language listing, give the name of the desired source file. For example, the command

```
cc –S add.c
```

creates an assembly language listing file named *add.s* and the command

```
cc –L mult.c
```

creates a listing file named *mult.L*. Note that both the –S and –L commands suppress subsequent compilation of the source file; they imply the –c option.

Thus, no program file is created and no linking is performed.

Another use of the –S option is to create an assembly language source file that may be optimized by hand and submitted to the XENIX assembler **as**. Although this method can be useful, optimizing should be left to the compiler whenever possible.

The –S and –L options apply to source files only; the compiler cannot create an assembly language listing file from an existing object file. Furthermore, the option in the command line must appear before the names of the files for which the assembly listing is to be saved.

2.7.2 Profiling a Program

You can examine the flow of execution of a program by adding "profiling" code to the program with the –p option. The profiling code automatically keeps a record of the number of times program functions are called during execution of the program. This record is written to the *mon.out* file and can be examined with the **prof** command.

For example, the command

```
cc –p main.c
```

adds profiling code to the program created from the source file *main.c*. The profiling code automatically calls the *monitor* function, which creates the *mon.out* file at normal termination of the program. The **prof** command and *monitor* function are described in detail in *prof*(CP) and *monitor*(S) in the XENIX *Reference Manual*.

The –p option must be given in any command line that references object files that contain profiling code. For example, if the command

```
cc –c –p f1.c f2.c
```

was used to create the object files *f1.o* and *f2.o*, then the command

```
cc –p f1.o f2.o
```

must be used to create an executable program from these files.

2.8 Controlling the C Preprocessor

The cc command provides a number of options that let you control the operation of the C preprocessor. These options let you define macros, create new search paths for include files, and suppress subsequent compilation of the source file.

2.8.1 Defining a Macro

You can define the value or meaning of a macro used in a source file by using the **–D** (for "define") option. The option lets you assign a value to a macro when you invoke the compiler, and is useful if you have used **if**, **ifdef**, and **ifndef** directives in your source files.

The option has the form

–D*name*[=*string*]

where *name* is the name of the macro and *string* is its value or meaning. If no *string* is given, the macro is assumed to be defined and its value is set to 1. For example, the command

```
cc –DNEED=2 main.c
```

sets the macro "NEED" to the value "2". This is the same as having the directive

```
#define NEED 2
```

in the source file. The command compiles the source file *main.c*, replacing every occurrence of "NEED" with "2".

The **–D** option is especially useful with the **ifdef** directive. You can use the option to determine which statements in the source are to be compiled. For example, suppose a source file, *main.c*, contains the directive

```
#ifdef NEED
```

but no explicit **define** directive for the macro "NEED". Then all statements following the **ifdef** directive are compiled only if you supply an explicit definition of "NEED" by using **–D**. For example, the command

```
cc –DNEED main.c
```

is sufficient to compile all statements following the **ifdef** directive, while the command

```
cc main.c
```

causes all those statements to be ignored.

You may use **–D** to define up to 20 macros on a command line. However, you cannot redefine a macro once it has been defined. If a file uses a macro, you must place the **–D** option before that file's name on the command line. For example, in the command

cc main.c –DNEED add.c

the macro "NEED" is defined for *add.c* but not defined for *main.c*.

2.8.2 Defining Include Directories

You can explicitly define the directories containing "include" files by using the –I (for "include") option. This option adds the given directory to a list of directories to be searched for include files. The directories in the list are searched whenever an include directive is encountered in the source file. The option has the form

–I*directoryname*

where *directoryname* is a valid pathname to a directory containing include files. For example, the command

cc –I/usr/joe/include main.c

causes the compiler to search the directory */usr/joe/include* for include files requested by the source file *main.c*.

The directories are searched in the order they are listed and only until the given include file is found. The */usr/include* directory is the default include directory and is always searched after directories given with –I.

2.8.3 Ignoring the Default Include Directories

You can prevent the C preprocessor from searching the default include directories by using the –X option. This option is generally used with the –I option to define the location of include files that have the same names as those found in the default directories, but which contain different definitions. For example, the command

cc –X –I/usr/joe/include main.c add.c

causes cc to look for all include files only in the directory */usr/joe/include*.

2.8.4 Saving a Preprocessed Source File

You can save a copy of the preprocessed source file by using the –P and –E options. The file is identical to the original source file except that all C preprocesor directives have been expanded or replaced. The –P option copies the result to the file named *filename*.i, where *filename* is the same name as the source file without the ".c" extension. The –E option copies the result to the standard output, and places a **#line** directive at the beginning and end of this output. You can save this output by redirecting it.

For example, the command

```
cc –P main.c
```

creates a preprocessed file *main.i* from the source file *main.c*, and the command

```
cc –E add.c >add.i
```

creates a preprocessed file from the source file *add.c*. The output is redirected to the file *add.i*.

Note that –P and –E suppress compilation of the source file. Thus, no object file or program is created.

2.9 Error Messages

The C compiler generates a broad range of error and warning messages to help you locate errors and potential problems in programs. In addition to compiler messages, the cc command also displays error messages generated by the XENIX C preprocessor and the XENIX assembler and linker programs. The following sections describe the form and meaning of the compiler error messages and warning messages you can encounter while using the cc command.

2.9.1 C Compiler Messages

The C compiler displays messages about syntactical and semantic errors in a source file, such as misplaced punctuation, Illegal use of operators, and undeclared variables. It also displays warning messages about statements containing potential problems caused by data conversions or the mismatch of types. Error and warning messages have the form

filename (*linenumber*) : *message*

where *filename* is the name of the source file being compiled, *linenumber* is the number of the line in the source file containing the error, and *message* is a self-explanatory description of the error or warning.

If an error is severe, the compiler displays a message and terminates the compilation. Otherwise, the compiler continues looking for other errors, but does not create an object file. If only warning messages are displayed, the compiler completes compilation and creates an object file.

You can avoid many C compiler errors by using the XENIX C program checker **lint** before compiling your C source files. **Lint** performs detailed error checking on a source file, and provides a list of actual errors and possible problems which may affect execution of the program. For a description of **lint**, see Chapter 3, "Lint: A C Program Checker".

2.9.2 Setting the Level of Warnings

You can set the level of warning messages produced by the compiler by using the –W option. This option directs the compiler to display messages about statements which may not be compiled as the programmer intends. Warnings indicate potential problems rather than actual errors. The option has the form

–W *number*

where *number* is a number in the range 0 to 3 giving the level of warnings. The levels are

Level	Warning
0	Suppresses all warning messages. Only messages about actual syntactical or semantic errors are displayed.
1	Warns about potentially missing statements, non-reachable statements, and other structural problems. Also, warn about overt type mismatches.
2	Warns about all type mismatches (strong typing).
3	Warns on all automatic data conversions.

If the option is not used, the default is level 1.

The higher option levels are especially useful in the earlier stages of program development when messages about potential problems are most helpful. The lower levels are best for compiling programs whose questionable statements are intentionally designed. For example, the command

```
cc –W 3 main.c
```

directs the compiler to perform the highest level of checking, and produces the greatest number warning messages. The command

```
cc –W 0 main.c
```

produces no warning messages. Note that the **–w** option has the same effect as **–W0**.

2.10 Using Advanced Options

The cc command provides a number of advanced programming options that give greater control over the compilation process and the final form of the executable program. The following sections describe a number of these options.

2.10.1 Creating Programs From Assembly Language Source Files

You can use the **cc** command to create executable programs from a combination of C source files and 8086/286 assembly language source files. Assembly language source files must contain 8086/286 instructions, as described in Chapter 7, "As: An Assembler," and must have the extension ".s".

When assembly language source files are given, the **cc** command invokes the XENIX assembler, **as**, to assemble the instructions and create an object file. The object file can then be linked with object files created by the compiler. For example, the command

```
cc main.c add.s
```

compiles the C source file *main.c*, but assembles the assemble language source file *add.s*. The resulting object files, *main.o* and *add.o*, are linked to form a single program.

When using assembly language routines with C programs, you must be sure to provide the correct interface for calls to and from C language functions. C functions require a specific calling and return sequence. Assembly language functions which fail to provide this interface will cause errors. See Appendix A, "Assembly Language Interface," in the XENIX *Library Guide*.

2.10.2 Using the near and far Keywords

The **near** and **far** keywords are special type modifiers that define the length and meaning of the address of a given variable. The **near** keyword defines an object with a 16 bit address. The **far** keyword defines an object with a full 32 bit segmented address. Any data item or function can be accessed.

The **near** and **far** keywords override the normal address length generated by the compiler for variables and functions. In small model programs, **far** lets you access data and functions in segments outside of the program. In middle and large model programs, **near** lets you access data with just an offset.

The examples in the following table illustrate the **far** and **near** keywords as used in declarations in a small model program.

Uses of near and far Keywords

Declaration	Address Size	Item Size
char c;	near (16 bits)	8 bits (data)
char far d;	far (32 bits)	8 bits (data)
char *p;	near (16 bits)	16 bits (near pointer)
char far *q;	near (16 bits)	32 bits (far pointer)
char * far r;	far (32 bits)	16 bits (near pointer)[1]
char far * far s;	far (32 bits)	32 bits (far pointer)[2]
int foo();	near (16 bits)	function returning 16 bits
int far foo();	far (32 bits)	function returning 16 bits[3]

Notes:

[1] This example has no meaning; it is shown for syntactic completeness only.

[2] This is similar to accessing data in a long model program.

[3] This example leads to trouble in most environments. The far call changes the CS register, and makes run time support unavailable.

The following example is from a middle model compilation:

```
int near foo();
```

This does a near call in an otherwise far (calling) program.

Since there is no type checking between items in separate source files, the **near** and **far** keywords should be used with great care.

2.10.3 Changing Word Order in Programs

The Microsoft C compiler automatically uses the standard 8086/286 word order for **long** type values. This order may cause problems when reading data files from programs created by other C compilers. You can change the word order for a given program by using the **–Mb** configuration option. This option causes the compiler to generate all long values in reverse word order, making the program compatible with programs created by other XENIX compilers.

Note that there are other portability issues which must be considered when creating C programs intended for several different XENIX systems. For an explanation of these issues, see Appendix B, "C Language Portability," in this guide.

2.10.4 Setting the Stack Size

You can set the size of the program stack by using the –F option. This option has the form

–F *num*

where *num* is the size (in bytes) of the program stack. The program stack is used for storage of function parameters and automatic variables. If the option is not used, a default stack size (usually 2 Kbytes) is set.

You can determine the stack requirements of a given program by using the **stackuse** program. This program analyzes C source files and computes the minimum stack requirement for all functions in the program. The program displays a warning if recursive functions are encountered; stack use requirements for recursive functions must be determined by the programmer. The **stackuse** program is described in *stackuse*(CP) in the XENIX *Reference Manual*.

Note that all programs created by cc have fixed stacks. This means the stack size cannot be increased during execution of the program. Therefore, a sufficient stack size must be given when compiling the program.

2.10.5 Using Modules, Segments, and Groups

"Module" is another name for the object file created by the C compiler. Every module has a name, and the cc command uses this name in error messages if problems are encountered during linking. The module name is usually the same as the source file's name (without the ".c" or ".s" extension). You can change this name using the –NM option. The option has the form

–NM *name*

where *name* can be any combination of letters and digits.

Changing a module's name is useful if the source file to be compiled is actually the output of a program preprocessor and generator, such as **lex** or **yacc**.

A "segment" is a contiguous block of binary code produced by the C compiler. Every module has two segments: a text segment containing the program instructions, and a data segment containing the program data. Each segment in every module has a name. This name is used by cc to define the order in which the segments of the program will appear in memory when loaded for execution. Text segments having the same name are loaded as a contiguous block of code. Data segments of the same name are also loaded as contiguous blocks.

Text and data segment names are normally created by the C compiler. These default names depend on the memory model chosen for the program. For example, in small model programs the text segment is named "_TEXT" and the data segment is named "_DATA". These names are the same for all small model modules, so all segments from all modules of a small model program are loaded as a contiguous block. In middle model programs, each text segment has a different name. In large model programs, each text and data segment has a different name. The default text and data segment names for middle and large model programs are given in the section "Segment and Module Names" given at the end of this chapter.

You can override the default names used by the C compiler (and override the default loading order) by using the **–NT** and **–ND** options. These options set the names of the text and data segments, in each module being compiled, to a given name. The options have the form

–NT *name*

and

–ND *name*

where *name* is any combination of letters and digits. These options are useful in middle and large model programs where there is no specific loading order. In these programs, you can guarantee contiguous loading for two or more segments by giving them the same name.

All text and data segments, whether or not they are loaded as contiguous blocks, are eventually loaded into one or more physical segments of memory. All segments in a physical segment are collectively called a "group".

All programs have at least two groups: a text group and a data group. Each group has a name. The text group is named "IGROUP" and the data group is named "DGROUP". The C compiler automatically applies these names to the text and data segments in each module. Thus, when the modules are eventually linked, all text segments belong to the same group, and all data segments belong to the same group.

Since a group corresponds to one physical segment, programs having more than 64 Kbytes each of text or data must be directed to two or more groups. (The limit per physical segment is 64 Kbytes.) You can create new groups by using the **–NGT** and **–NGD** options. The options have the form

–NGT *name*

and

–NGD *name*

where *name* is any combination of letters and digits. These options set the

group name for the text and data segments of the given modules. All segments with the same name are loaded in the same physical segment. There will be one physical segment for each group named.

2.11 Compiler Summary

The following sections summarize cc options and memory models.

2.11.1 Cc Options

The following is a complete list of cc options:

-c Creates a linkable object file for each source file.

-C Preserves comments when preprocessing a file (only when -P or -E).

-D name [= string]
Defines *name* to the preprocessor. The value is *string* or 1.

-E Preprocesses each source file, copying the result to the standard output.

-F num
Sets the size of the program stack.

-i Creates separate instruction and data spaces for small model programs.

-I pathname
Adds *pathname* to the list of directories to be searched for #include files.

-K Removes stack probes from a program.

-l name
Search library *name* for unresolved function names.

-L Creates an assembler listing file containing assembled code and assembly source instructions.

-M string
Sets the program configuration. The *string* may be any combination of "s" (small model), "m" (middle model), "l" (large model), "e" (enable far and near keywords), "2" (enables 286 code generation), "b" (reverse word order), and "t" (sets data threshold for largest item in a segment). The "s", "m", and "l" are mutually exclusive.

-nl num
Sets the maximum length of external symbols.

-ND name
Sets the data segment name.

-NGD name
Sets the data group name.

-NGT name
Sets the text group name.

-NM name
Sets the module name.

-NT name
Sets the text segment name.

-o filename
Makes *filename* the name of the final executable program.

-O Invokes the object code optimizer.

-p Adds code for program profiling.

-P Preprocesses source files and sends output to files with the extension ".i".

-S Creates an assembly source listing.

-V string
Copies *string* to the object file.

-w Suppresses compiler warning messages.

-W num
Sets the output level for compiler warning messages.

-X Removes the standard directories from the list of directories to be searched for #include files.

2.11.2 Memory Models

The following table defines the number of text and data segments for the four different program memory models. This table also lists the segment register values.

Model	Text	Data	Segment Registers
Small	1*	1*	CS=DS=SS
Middle	1 per module	1	DS=SS
Large	1 per module	1 per module	

* -- In impure-text small module programs, text and data occupy the same segment. In pure-text programs, they occupy different segments.

2.11.3 Pointer and Integer Sizes

The following table defines the sizes (in bits) of integers (**int** type), and text and data pointers, in each program memory model.

Model	Data Pointer	Text Pointer	Integer
Small	16	16	16
Middle	16	32	16
Large	32	32	16

2.11.4 Segment and Module Names

The following table lists the default text and data segment names, and the default module name, for each object file.

Model	Text	Data	Module
Small	_TEXT	_DATA	*filename*
Middle	*module*_TEXT	_DATA	*filename*
Large	*module*_TEXT	*module*_DATA	*filename*

Chapter 3
Lint: A C Program Checker

3.1 Introduction 3-1

3.2 Invoking *lint* 3-1

3.3 Checking for Unused Variables and Functions 3-2

3.4 Checking Local Variables 3-3

3.5 Checking for Unreachable Statements 3-4

3.6 Checking for Infinite Loops 3-4

3.7 Checking Function Return Values 3-5

3.8 Checking for Unused Return Values 3-6

3.9 Checking Types 3-6

3.10 Checking Type Casts 3-7

3.11 Checking for Nonportable Character Use 3-7

3.12 Checking for Assignment of longs to ints 3-8

3.13 Checking for Strange Constructions 3-8

3.14 Checking for Use of Older C Syntax 3-9

3.15 Checking Pointer Alignment 3-10

3.16 Checking Expression Evaluation Order 3-10

3.17 Embedding Directives 3-11

3.18 Checking For Library Compatibility 3-12

3.1 Introduction

This chapter explains how to use the C program checker **lint**. The program examines C source files and warns of errors or misconstructions that may cause errors during compilation of the file or during execution of the compiled file.

In particular, **lint** checks for

Unused functions and variables

Unknown values in local variables

Unreachable statements and infinite loops

Unused and misused return values

Inconsistent types and type casts

Mismatched types in assignments

Nonportable and old fashioned syntax

Strange constructions

Inconsistent pointer alignment and expression evaluation order

The **lint** program and the C compiler are generally used together to check and compile C language programs. Although the C compiler rapidly and efficiently compiles C language source files, it does not perform the sophisticated type and error checking required by many programs. The **lint** program, on the other hand, provides thorough checking of source files without compiling.

3.2 Invoking *lint*

You can invoke **lint** by typing its name at the shell command line. The command has the form

lint [*option*] ... *filename* ... *lib* ...

where *option* is a command option that defines how the checker should operate, *filename* is the name of the C language source file to be checked, and *lib* is the name of a library to check. You can give more than one option, filename , or library name in the command as long as you use spaces to separate them. If you give two or more filenames, **lint** assumes that the files form a complete program and checks the files accordingly. For example, the command

lint main.c add.c

treats *main.c* and *add.c* as two parts of a complete program.

If **lint** discovers errors or inconsistencies in a source file, it produces messages describing the problem. The message has the form

filename (*num*): *description*

where *filename* is the name of the source file containing the problem, *num* is the number of the line in the source containing the problem, and *description* is a description of the problem. For example, the message

main.c (3): warning: x unused in function main

shows that the variable "x", defined in line three of the source file *main.c*, is not used anywhere in the file.

3.3 Checking for Unused Variables and Functions

The **lint** program checks for unused variables and functions by seeing if each declared variable and function is used at least once in the source file. The program considers a variable or function used if the name appears in at least one statement. It is not considered used if it only appears on the left side of an assignment. For example, in the following program fragment

```
main ()
{
        int x,y,z;

        x=1; y=2; z=x+y;
```

the variables "x" and "y" are considered used, but variable "z" is not.

Unused variables and functions often occur during the development of large programs. It is not uncommon for a programmer to remove all references to a variable or function from a source file, but forget to remove its declaration. Such unused variables and functions rarely cause working programs to fail, but do make programs harder to understand and change. Checking for unused variables and functions can also help you find variables or functions that you intended to used but accidentally have left out of the program.

Note that the **lint** program does not report a variable or function unused if it is explicitly declared with the **extern** storage class. Such a variable or function is assumed to be used in another source file.

You can direct **lint** to ignore all the external declarations in a source file by using the –x (for "external") option. This option causes the program checker to skip any line that begins with the **extern** storage class. The –x option is typically used to save time when checking a program, especially if all external declarations are known to be valid.

Some programming styles require functions that perform closely related tasks to have the same number and type of arguments, regardless of whether these arguments are used. Under normal operation, **lint** reports any argument not used as an unused variable, but you can direct **lint** to ignore unused arguments by using the –v option. The –v option causes **lint** to ignore all unused function arguments except for those declared with **register** storage class. The program considers unused arguments of this class to be a preventable waste of the register resources of the computer.

You can direct **lint** to ignore all unused variables and functions by using the –**u** (for "unused") option. This option prevents **lint** from reporting variables and functions it considers unused. The –**u** option is typically used when checking a source file that contains just a portion of a large program. Such source files usually contain declarations of variables and functions that are intended to be used in other source files and are not explicitly used within the file. Since **lint** can only check the given file, it assumes that such variables or functions are unused and and reports them as errors whenever the –**u** option is not given.

3.4 Checking Local Variables

The **lint** program checks all local variables to ensure that they are set to a value before being used. Since local variables have either automatic or register storage class, their values at the start of the program or function cannot be known. Using such a variable before assigning a value to it is an error.

The **lint** program checks the local variables by searching for the first assignment in which the variable receives a value, and for the first statement or expression in which the variable is used. If the first assignment appears later than the first use, **lint** considers the variable inappropriately used. For example, in the program fragment

```
char c;

if ( c != EOT )
        c = getchar();
```

lint warns that the the variable "c" is used before it is assigned.

If a variable is used in the same statement in which it is assigned for the first time, **lint** determines the order of evaluation of the statement and displays an appropriate message. For example, in the program fragment

```
int i,total;

scanf("%d", &i);
total = total + i;
```

lint warns that the variable "total" is used before it is set, since it appears on the right side of the same statement that assigns its first value.

Static and external variables are always initialized to zero before program execution begins, so **lint** does not report such variables if they are used before being set to a value.

3.5 Checking for Unreachable Statements

The **lint** program checks for unreachable statements. Unreachable statements are unlabeled statements that immediately follow a **goto**, **break**, **continue**, or **return** statement. During execution of a program, the unreachable statements never receive execution control and therefore are considered wasteful. For example, in the program fragment

```
int x,y;

return (x+y);
exit (1);
```

the function call *exit* after the return statement is unreachable.

Unreachable statements are common when developing programs containing large case constructions, or loops containing break and continue statements. Such statements are wasteful and should be removed when convenient.

During normal operation, **lint** reports all unreachable break statements. Unreachable break statements are relatively common (some programs created by the *yacc* and *lex* programs contain hundreds), so it may be desirable to suppress these reports. You can direct **lint** to suppress the reports by using the **–b** option.

Note that **lint** assumes that all functions eventually return control, so it does not report as unreachable any statement that follows a function that takes control and never returns it. For example, in the program fragment

```
exit (1);
return;
```

the call to *exit* causes the **return** statement to become an unreachable statement, but **lint** does not report it as such.

3.6 Checking for Infinite Loops

The **lint** program checks for infinite loops and for loops which are never executed. For example, the statements

```
while (1) { }
```

and

```
for (;;) {}
```

are both considered infinite loops. While the statements

```
while (0) { }
```

and

```
for (0;0;) { }
```

are never executed.

Although some valid programs have such loops, they are generally considered errors.

3.7 Checking Function Return Values

The **lint** program checks to ensure that a function returns a meaningful value if a return value is expected. Some functions return values that are never used; some programs incorrectly use function values that have never been returned. **Lint** addresses these problems in a number of ways.

Within a function definition, the appearance of both

```
return (expr);
```

and

```
return ;
```

statements is cause for alarm. In this case, **lint** produces the following error message:

```
function name contains return(e) and return
```

It is difficult to detect when a function return is implied by the flow of control reaching the end of the given function. This is demonstrated with a simple example.

```
f (a)
{
        if (a)
                return (3);
        g ();
}
```

In this example, if the variable "a" is false, then *f* will call the function *g* and then return with no defined return value. This will trigger a report from **lint**. If *g*, like *exit*, never returns, the message will still be produced when in fact

nothing is wrong. In practice, potentially serious bugs can be discovered with this feature. It also accounts for a substantial fraction of the undeserved error messages produced by **lint**.

3.8 Checking for Unused Return Values

The **lint** program checks for cases where a function returns a value, but the value is rarely if ever used. **Lint** considers functions that return unused values to be inefficient, and functions that return rarely used values to be a result of bad programming style.

Lint also checks for cases where a function does not return a value but the value is used anyway. This is considered a serious error.

3.9 Checking Types

Lint enforces the type checking rules of C more strictly than the C compiler. The additional checking occurs in four major areas.

1. Across certain binary operators and implied assignments
2. At the structure selection operators
3. Between the definition and uses of functions
4. In the use of enumerations

There are a number of operators that have an implied balancing between types of operands. The assignment, conditional, and relational operators have this property. The argument of a **return** statement, and expressions used in initialization also suffer similar conversions. In these operations, **char**, **short**, **int**, **long**, **unsigned**, **float**, and **double** types may be freely intermixed. The types of pointers must agree exactly, except that arrays of x's can be intermixed with pointers to x's.

The type checking rules also require that, in structure references, the left operand of a pointer arrow symbol (->) be a pointer to a structure, the left operand of a period (.) be a structure, and the right operand of these operators be a member of the structure implied by the left operand. Similar checking is done for references to unions.

Strict rules apply to function argument and return value matching. The types **float** and **double** may be freely matched, as may the types **char**, **short**, **int**, and **unsigned**. Pointers can also be matched with the associated arrays. Aside from these relaxations in type checking, all actual arguments must agree in type with their declared counterparts.

Lint checks to ensure are made that enumeration variables or members are not mixed with other types or other enumerations. It also ensures that the only operations applied to enumerated variables are assignment (=), initialization, equals (==), and not-equals (!=). Enumerations may also be function arguments and return values.

3.10 Checking Type Casts

The type cast feature in C was introduced largely as an aid to producing more portable programs. Consider the assignment

```
p = 1 ;
```

where "p" is a character pointer. Lint reports this as suspect. But consider the assignment

```
p = (char *)1 ;
```

in which a cast has been used to convert the integer to a character pointer. The programmer obviously had a strong motivation for doing this, and has clearly signaled his intentions. On the other hand, if this code is moved to another machine, it should be looked at carefully. The –c option controls the printing of comments about casts. When –c is in effect, casts are not checked, and all legal casts are passed without comment, no matter how strange the type mixing seems to be.

3.11 Checking for Nonportable Character Use

Lint flags certain comparisons and assignments as illegal or nonportable. For example, the fragment

```
char c;
.
.
.
if( (c = getchar()) < 0 ) ...
```

works on some machines, but fails on machines where characters always take on positive values. In this case, lint issues the message:

```
nonportable character comparison
```

The solution is to declare "c" an integer, since *getchar* is actually returning integer values.

A similar issue arises with bitfields. When assignments of constant values are made to bitfields, the field may be too small to hold the value. This is especially true where on some machines bitfields are considered as signed quantities.

Although a 2-bit field with **int** type cannot hold the value 3, a 2-bit field with **unsigned** type can.

3.12 Checking for Assignment of longs to ints

Problems may arise from the assignment of **long** values to an **int** values, because of a loss in accuracy in the assignment. This may happen in programs that have been incompletely converted by changing type definitions with **typedef**. When a **typedef** variable is changed from **int** to **long**, the program can stop working because some intermediate results may be assigned to integer values, losing accuracy. Since there are a number of legitimate reasons for assigning longs to integers, you may wish to suppress detection of these assignments by using the **–a** option.

3.13 Checking for Strange Constructions

Several perfectly legal, but somewhat strange, constructions are flagged by **lint**. The generated messages encourage better code quality, clearer style, and may even point out bugs. For example, in the statement

```
*p++ ;
```

the star (*) does nothing, so **lint** prints

```
null effect
```

The program fragment

```
unsigned x ;
if (x < 0) ...
```

is also strange since the test will never succeed.

Similarly, the test

```
if (x > 0)
```

is equivalent to

```
if (x != 0)
```

which may not be the intended action. In these cases, **lint** prints the message

```
degenerate unsigned comparison
```

If you use

```
if( 1 != 0 ) ...
```

then **lint** reports

```
constant in conditional context
```

since the comparison of 1 with 0 gives a constant result.

Another construction detected by **lint** involves operator precedence. Bugs that arise from misunderstandings about the precedence of operators can be accentuated by spacing and formatting, making such bugs extremely hard to find. For example, the statements

```
if( x&077 == 0 ) ...
```

or

```
x<<2 + 40
```

probably do not do what is intended. The best solution is to place parentheses around such expressions. **Lint** encourages this by printing an appropriate message.

Finally, **lint** checks variables that are redeclared in inner blocks in a way that conflicts with their use in outer blocks. This is legal, but is considered bad style, usually unnecessary, and frequently a bug.

If you do not wish these heuristic checks, you can suppress them by using the **–h** option.

3.14 Checking for Use of Older C Syntax

Lint checks for older C constructions. These fall into two classes: assignment operators and initialization.

The older forms of assignment operators (e.g., =+, =–, ...) can cause ambiguous expressions, such as

```
a =-1 ;
```

which could be taken as either

```
a =- 1 ;
```

or

```
a = -1 ;
```

The situation is especially perplexing if this kind of ambiguity arises as the

result of a macro substitution. The newer, and preferred operators (e.g., +=, -=) have no such ambiguities. To encourage the abandonment of the older forms, **lint** checks for occurrences of these old-fashioned operators.

A similar issue arises with initialization. The older language allowed

```
int x 1;
```

to initialize "x" to 1. This causes syntactic difficulties. For example

```
int x (-1);
```

looks somewhat like the beginning of a function declaration

```
int x (y) { ...
```

and the compiler must read past "x" to determine what the declaration really is. The problem is even more perplexing when the initializer involves a macro. The current C syntax places an equal sign between the variable and the initializer

```
int x = -1;
```

This form is free of any possible syntactic ambiguity.

3.15 Checking Pointer Alignment

Certain pointer assignments may be reasonable on some machines, and illegal on others, due to alignment restrictions. For example, on some machines it is reasonable to assign integer pointers to double pointers, since double precision values may begin on any integer boundary. On other machines, however, double precision values must begin on even word boundaries; thus, not all such assignments make sense. **Lint** tries to detect cases where pointers are assigned to other pointers, and such alignment problems might arise. The message

```
possible pointer alignment problem
```

results from this situation.

3.16 Checking Expression Evaluation Order

In complicated expressions, the best order in which to evaluate subexpressions may be highly machine-dependent. For example, on machines in which the stack runs backwards, function arguments will probably best be evaluated from right to left; on machines with a stack running forward, left to right is probably best. Function calls embedded as arguments of other functions may or may not be treated in the same way as ordinary arguments. Similar issues arise with other operators that have side effects, such as the assignment

operators and the increment and decrement operators.

To ensure maximum efficiency of C on a particular machine, the C language leaves the order of evaluation of complicated expressions up to the compiler. Various C compilers have considerable differences in the order in which they will evaluate complicated expressions. In particular, if any variable is changed by a side effect, and also used elsewhere in the same expression, the result is undefined.

Lint checks for the important special case where a simple scalar variable is affected. For example, the statement

```
a[i] = b[i++] ;
```

will draw the comment

```
warning: i evaluation order undefined
```

3.17 Embedding Directives

There are occasions when the programmer is smarter than **lint**. There may be valid reasons for illegal type casts, functions with a variable number of arguments, and other constructions that **lint** finds objectionable. Moreover, as specified in the above sections, the flow of control information produced by **lint** often has blind spots, causing occasional spurious messages about perfectly reasonable programs. Some way of communicating with **lint**, typically to turn off its output, is desirable. Therefore, a number of words are recognized by **lint** when they are embedded in comments in a C source file. These words are called directives. **Lint** directives are invisible to the compiler.

The first directive discussed concerns flow of control information. If a particular place in the program cannot be reached, this can be asserted at the appropriate spot in the program with the directive

```
/* NOTREACHED */
```

Similarly, if you desire to turn off strict type checking for the next expression, use the directive

```
/* NOSTRICT */
```

The situation reverts to the previous default after the next expression. The –v option can be turned on for one function with the directive

```
/* ARGSUSED */
```

Comments about a variable number of arguments in calls to a function can be turned off by preceding the function definition with the directive

```
/* VARARGS */
```

In some cases, it is desirable to check the first several arguments, and leave the later arguments unchecked. You can define the number of arguments to be checked by placing a digit (giving this number) immediately after the VARARGS keyword. For example,

```
/* VARARGS2 */
```

causes only the first two arguments to be checked. Finally, the directive

```
/* LINTLIBRARY */
```

at the head of a file identifies this file as a library declaration file, which is discussed in the next section.

3.18 Checking For Library Compatibility

Lint accepts certain library directives, such as

```
-ly
```

and tests the source files for compatibility with these libraries. This testing is done by accessing library description files whose names are constructed from the library directives. These files all begin with the directive

```
/* LINTLIBRARY */
```

which is followed by a series of dummy function definitions. These definitions indicate whether a function returns a value, what type a function's return type is, and the number and types of arguments expected by the function. The "VARARGS" and "ARGSUSED" directives can be used to specify features of the library functions.

Lint library files are processed almost exactly like ordinary source files. The only difference is that functions that are defined in a library file, but are not used in a source file, draw no comments. **Lint** does not simulate a full library search algorithm, and checks to see if the source files contain redefinitions of library routines.

By default, **lint** checks the programs it is given against a standard library file, which contains descriptions of the programs that are normally loaded when a C program is run. When the **-p** option is in effect, the portable library file is checked. This library contains descriptions of the standard I/O library routines which are expected to be portable across various machines. The **-n** option can be used to suppress all library checking.

Chapter 4
Make: A Program Maintainer

4.1 Introduction 4-1

4.2 Creating a Makefile 4-1

4.3 Invoking Make 4-3

4.4 Using Pseudo-Target Names 4-4

4.5 Using Macros 4-5

4.6 Using Shell Environment Variables 4-8

4.7 Using the Built-In Rules 4-9

4.8 Changing the Built-in Rules 4-10

4.9 Using Libraries 4-12

4.10 Troubleshooting 4-13

4.11 Using Make: An Example 4-13

4.1 Introduction

The **make** program provides an easy way to automate the creation of large programs. **Make** reads commands from a user-defined "makefile" that lists the files to be created, the commands that create them, and the files from which they are created. When you direct **make** to create a program, it verifies that each file on which the program depends is up to date, then creates the program by executing the given commands. If a file is not up to date, **make** updates it before creating the program. **Make** updates a program by executing explicitly given commands, or one of the many built-in commands.

This chapter explains how to use **make** to automate medium-sized programming projects. It explains how to create makefiles for each project, and how to invoke **make** for creating programs and updating files. For more details about the program, see *make* (CP) in the XENIX *Reference Manual*.

4.2 Creating a Makefile

A makefile contains one or more lines of text called dependency lines. A dependency line shows how a given file depends on other files and what commands are required to bring a file up to date. A dependency line has the form

target ... : [*dependent* ...] [; *command* ...]

where *target* is the filename of the file to be updated, *dependent* is the filename of the file on which the target depends, and *command* is the XENIX command needed to create the target file. Each dependency line must have at least one command associated with it, even if it is only the null command (;).

You may give more than one target filename or dependent filename if desired. Each filename must be separated from the next by at least one space. The target filenames must be separated from the dependent filenames by a colon (:). Filenames must be spelled as defined by the XENIX system. Shell metacharacters, such as star (*) and question mark (?), can also be used.

You may give a sequence of commands on the same line as the target and dependent filenames, if you precede each command with a semicolon (;). You can give additional commands on following lines by beginning each line with a tab character. Commands must be given exactly as they would appear on a shell command line. The at sign (@) may be placed in front of a command to prevent **make** from displaying the command before executing it. Shell commands, such as *cd*(C), must appear on single lines; they must not contain the backslash (\) and newline character combination.

You may add a comment to a makefile by starting the comment with a number sign (#) and ending it with a newline character. All characters after the number sign are ignored. Comments may be place at the end of a dependency

line if desired. If a command contains a number sign, it must be enclosed in double quotation marks (").

If a dependency line is too long, you can continue it by typing a backslash (\) and a newline character.

The makefile should be kept in the same directory as the given source files. For convenience, the filenames **makefile**, **Makefile**, *s.makefile*, and *s.Makefile* are provided as default filenames. These names are used by **make** if no explicit name is given at invocation. You may use one of these names for your makefile, or choose one of your own. If the filename begins with the *s.* prefix, **make** assumes that it is an SCCS file and invokes the appropriate SCCS command to retrieve the lastest version of the file.

To illustrate dependency lines, consider the following example. A program named *prog* is made by linking three object files, *x.o*, *y.o*, and *z.o*. These object files are created by compiling the C language source files *x.c*, *y.c*, and *z.c*. Furthermore, the files *x.c* and *y.c* contain the line

```
#include "defs"
```

This means that *prog* depends on the three object files, the object files depend on the C source files, and two of the source files depend on the include file *defs*. You can represent these relationships in a makefile with the following lines.

```
prog: x.o y.o z.o
        cc x.o y.o z.o -o prog
x.o: x.c defs
        cc -c x.c
y.o: y.c defs
        cc -c y.c
z.o: z.c
        cc -c z.c
```

In the first dependency line, *prog* is the target file and *x.o*, *y.o*, and *z.o* are its dependents. The command sequence

```
cc x.o y.o z.o -o prog
```

on the next line tells how to create *prog* if it is out of date. The program is out of date if any one of its dependents has been modified since *prog* was last created.

The second, third, and fourth dependency lines have the same form, with the *x.o*, *y.o*, and *z.o* files as targets and *x.c*, *y.c*, *z.c*, and *defs* files as dependents. Each dependency line has one command sequence which defines how to update the given target file.

4.3 Invoking Make

Once you have a makefile and wish to update and modify one or more target files in the file, you can invoke **make** by typing its name and optional arguments. The invocation has the form

make [*option*] ... [*macdef*] ... [*target*] ...

where *option* is a program option used to modify program operation, *macdef* is a macro definition used to give a macro a value or meaning, and *target* is the filename of the file to be updated. It must correspond to one of the target names in the makefile. All arguments are optional. If you give more than one argument, you must separate them with spaces.

You can direct **make** to update the first target file in the makefile by typing just the program name. In this case, **make** searches for the files **makefile**, **Makefile**, *s.makefile*, and *s.Makefile* in the current directory, and uses the first one it finds as the makefile. For example, assume that the current makefile contains the dependency lines given in the last section. Then the command

```
make
```

compares the current date of the *prog* program with the current date each of the object files *x.o*, *y.o*, and *z.o*. It recreates *prog* if any changes have been made to any object file since *prog* was last created. It also compares the current dates of the object files with the dates of the four source files *x.c*, *y.c*, *z.c*, or *defs*, and recreates the object files if the source files have changed. It does this before recreating *prog* so that the recreated object files can be used to recreate *prog*. If none of the source or object files have been altered since the last time *prog* was created, **make** announces this fact and stops. No files are changed.

You can direct **make** to update a given target file by giving the filename of the target. For example,

```
make x.o
```

causes **make** to recompile the *x.o* file, if the *x.c* or *defs* files have changed since the object file was last created. Similarly, the command

```
make x.o z.o
```

causes **make** to recompile *x.o* and *z.o* if the corresponding dependents have been modified. **Make** processes target names from the command line in a left to right order.

You can specify the name of the makefile you wish **make** to use by giving the **–f** option in the invocation. The option has the form

–f filename

where *filename* is the name of the makefile. You must supply a full pathname if the file is not in the current directory. For example, the command

make –f makeprog

reads the dependency lines of the makefile named **makeprog** found in the current directory. You can direct **make** to read dependency lines from the standard input by giving "-" as the *filename*. **Make** reads the standard input until the end-of-file character is encountered.

You may use the program options to modify the operation of the **make** program. The following list describes some of the options.

–p Prints the complete set of macro definitions and dependency lines in a makefile.

–i Ignores errors returned by XENIX commands.

–k Abandons work on the current entry, but continues on other branches that do not depend on that entry.

–s Executes commands without displaying them.

–r Ignores the built-in rules.

–n Displays commands but does not execute them. **Make** even displays lines beginning with the at sign (@).

–e Ignores any macro definitions that attempt to assign new values to the shell's environment variables.

–t Changes the modification date of each target file without recreating the files.

Note that **make** executes each command in the makefile by passing it to a separate invocation of a shell. Because of this, care must be taken with certain commands (e.g., **cd** and shell control commands) that have meaning only within a single shell process; the results are forgotten before the next line is executed. If an error occurs, **make** normally stops the command.

4.4 Using Pseudo-Target Names

It is often useful to include dependency lines that have pseudo-target names, i.e., names for which no files actually exist or are produced. Pseudo-target

names allow **make** to perform tasks not directly connected with the creation of a program, such as deleting old files or printing copies of source files. For example, the following dependency line removes old copies of the given object files when the pseudo-target name "cleanup" is given in the invocation of **make**.

```
cleanup :
        rm  x.o y.o z.o
```

Since no file exists for a given pseudo-target name, the target is always assumed to be out of date. Thus the associated command is always executed.

Make also has built-in pseudo-target names that modify its operation. The pseudo-target name ".IGNORE" causes **make** to ignore errors during execution of commands, allowing **make** to continue after an error. This is the same as the –i option. (*Make* also ignores errors for a given command if the command string begins with a hyphen (–).)

The pseudo-target name ".DEFAULT" defines the commands to be executed either when no built-in rule or user-defined dependency line exists for the given target. You may give any number of commands with this name. If ".DEFAULT" is not used and an undefined target is given, **make** prints a message and stops.

The pseudo-target name ".PRECIOUS" prevents dependents of the current target from being deleted when **make** is terminated using the INTERRUPT or QUIT key, and the pseudo-target name ".SILENT" has the same effect as the –s option.

4.5 Using Macros

An important feature of a makefile is that it can contain macros. A macro is a short name that represents a filename or command option. The macros can be defined when you invoke **make**, or in the makefile itself.

A macro definition is a line containing a name, an equal sign (=), and a value. The equal sign must not be preceded by a colon or a tab. The name (string of letters and digits) to the left of the equal sign (trailing blanks and tabs are stripped) is assigned the string of characters following the equal sign (leading blanks and tabs are stripped.) The following are valid macro definitions:

```
2 = xyz
abc = -ll -ly
LIBES =
```

The last definition assigns "LIBES" the null string. A macro that is never explicitly defined has the null string as its value.

A macro is invoked by preceding the macro name with a dollar sign; macro names longer than one character must be placed in parentheses. The name of the macro is either the single character after the dollar sign or a name inside parentheses. The following are valid macro invocations.

```
$(CFLAGS)
$2
$(xy)
$Z
$(Z)
```

The last two invocations are identical.

Macros are typically used as placeholders for values that may change from time to time. For example, the following makefile uses a macro for the names of object files to be link and one for the names of the library.

```
OBJECTS = x.o y.o z.o
LIBES = -lln
prog: $(OBJECTS)
        cc $(OBJECTS)  $(LIBES)  -o prog
```

If this makefile is invoked with the command

```
make
```

it will load the three object files with the *lex* library specified with the **-lln** option.

You may include a macro definition in a command line. A macro definition in a command line has the same form as a macro definition in a makefile. If spaces are to be used in the definition, double quotation marks must be used to enclose the definition. Macros in a command line override corresponding definitions found in the makefile. For example, the command

```
make  "LIBES=-lln -lm"
```

loads assigns the library options **-lln** and **-lm** to "LIBES".

You can modify all or part of the value generated from a macro invocation without changing the macro itself by using the "substitution sequence". The sequence has the form

name : *st1* =[*st2*]

where *name* is the name of the macro whose value is to be modified, *st1* is the character or characters to be modified, and *st2* is the character or characters to replace the modified characters. If *st2* is not given, *st1* is replaced by a null character.

The substitution sequence is typically used to allow user-defined metacharacters in a makefile. For example, suppose that ".x" is to be used as a metacharacter for a prefix and suppose that a makefile contains the definition

```
FILES = prog1.x prog2.x prog3.x
```

Then the macro invocation

```
$(FILES : .x=.o)
```

generates the value

```
prog1.o prog2.o prog3.o
```

The actual value of "FILES" remains unchanged.

Make has five built-in macros that can be used when writing dependency lines. The following is a list of these macros.

$*	Contains the name of the current target with the suffix removed. Thus if the current target is *prog.o*, $* contains *prog*. It may be used in dependency lines that redefine the built-in rules.
$@	Contains the full pathname of the current target. It may be used in dependency lines with user-defined target names.
$<	Contains the filename of the dependent that is more recent than the given target. It may be used in dependency lines with built-in target names or the **.DEFAULT** pseudo-target name.
$?	Contains the filenames of the dependents that are more recent than the given target. It may be used in dependency lines with user-defined target names.
$%	Contains the filename of a library member. It may be used with target library names (see the section "Using Libraries" later in this chapter). In this case, $@ contains the name of the library and $% contains the name of the library member.

You can change the meaning of a built-in macro by appending the **D** or **F** descriptor to its name. A built-in macro with the **D** descriptor contains the name of the directory containing the given file. If the file is in the current directory, the macro contains ".". A macro with the **F** descriptor contains the name of the given file with the directory name part removed. The **D** and **F** descriptor must not be used with the **$?** macro.

4.6 Using Shell Environment Variables

Make provides access to current values of the shell's environment variables such as "HOME", "PATH", and "LOGIN". Make automatically assigns the value of each shell variable in your environment to a macro of the same name. You can access a variable's value in the same way that you access the value of explicitly defined macros. For example, in the following dependency line, "$(HOME)" has the same value as the user's "HOME" variable.

```
prog :
        cc $(HOME)/x.o $(HOME)/y.o /usr/pub/z.o
```

Make assigns the shell variable values after it assigns values to the built-in macros, but before it assigns values to user-specified macros. Thus, you can override the value of a shell variable by explicitly assigning a value to the corresponding macro. For example, the following macro definition causes **make** to ignore the current value of the "HOME" variable and use */usr/pub* instead.

```
HOME = /usr/pub
```

If a makefile contains macro definitions that override the current values of the shell variables, you can direct **make** to ignore these definitions by using the **–e** option.

Make has two shell variables, "MAKE" and "MAKEFLAGS", that correspond to two special-purpose macros.

The "MAKE" macro provides a way to override the **–n** option and execute selected commands in a makefile. When "MAKE" is used in a command, **make** will always execute that command, even if **–n** has been given in the invocation. The variable may be set to any value or command sequence.

The "MAKEFLAGS" macro contains one or more **make** options, and can be used in invocations of **make** from within a makefile. You may assign any **make** options to "MAKEFLAGS" except **–f**, **–p**, and **–d**. If you do not assign a value to the macro, **make** automatically assigns the current options to it, i.e., the options given in the current invocation.

The "MAKE" and "MAKEFLAGS" variables, together with the **–n** option, are typically used to debug makefiles that generate entire software systems. For example, in the following makefile, setting "MAKE" to "make" and invoking this file with the **–n** options displays all the commands used to generate the programs *prog1*, *prog2*, and *prog3* without actually executing them.

```
system : prog1 prog2 prog3
        @echo System complete.

prog1 : prog1.c
        $(MAKE) $(MAKEFLAGS) prog1

prog2 : prog2.c
        $(MAKE) $(MAKEFLAGS) prog2

prog3 : prog3.c
        $(MAKE) $(MAKEFLAGS) prog3
```

4.7 Using the Built-In Rules

Make provides a set of built-in dependency lines, called built-in rules, that automatically check the targets and dependents given in a makefile, and create up-to-date versions of these files if necessary. The built-in rules are identical to user-defined dependency lines except that they use the suffix of the filename as the target or dependent instead of the filename itself. For example, **make** automatically assumes that all files with the suffix *.o* have dependent files with the suffixes *.c* and *.s*.

When no explicit dependency line for a given file is given in a makefile, **make** automatically checks the default dependents of the file. It then forms the name of the dependents by removing the suffix of the given file and appending the predefined dependent suffixes. If the given file is out of date with respect to these default dependents, **make** searches for a built-in rule that defines how to create an up-to-date version of the file, then executes it. There are built-in rules for the following files.

.o	Object file
.c	C source file
.r	Ratfor source file
.f	Fortran source file
.s	Assembler source file
.y	Yacc-C source grammar
.yr	Yacc-Ratfor source grammar
.l	Lex source grammar

For example, if the file *x.o* is needed and there is an *x.c* in the description or directory, it is compiled. If there is also an *x.l*, that grammar would be run through *lex* before compiling the result.

The built-in rules are designed to reduce the size of your makefiles. They provide the rules for creating common files from typical dependents. Reconsider the example given in the section "Creating a Makefile". In this example, the program *prog* depended on three object files *x.o*, *y.o*, and *z.o*. These files in turn depended on the C language source files *x.c*, *y.c*, and *z.c*.

The files *x.c* and *y.c* also depended on the include file *defs*. In the original example each dependency and corresponding command sequence was explicitly given. Many of these dependency lines were unnecessary, since the built-in rules could have been used instead. The following is all that is needed to show the relationships between these files.

```
prog: x.o y.o z.o
        cc x.o y.o z.o -o prog

x.o y.o: defs
```

In this makefile, *prog* depends on three object files, and an explicit command is given showing how to update *prog*. However, the second line merely shows that two objects files depend on the include file *defs*. No explicit command sequence is given on how to update these files if necessary. Instead, **make** uses the built-in rules to locate the desired C source files, compile these files, and create the necessary object files.

4.8 Changing the Built-in Rules

You can change the built-in rules by redefining the macros used in these lines or by redefining the commands associated with the rules. You can display a complete list of the built-in rules and the macros used in the rules by typing

```
make -fp - 2>/dev/null </dev/null
```

The rules and macros are displayed at the standard output.

The macros of the built-in dependency lines define the names and options of the compilers, program generators, and other programs invoked by the built-in commands. Make automatically assigns a default value to these macros when you start the program. You can change the values by redefining the macro in your makefile. For example, the following built-in rule contains three macros, "CC", "CFLAGS", and "LOADLIBES".

```
.c :
        $(CC) $(CFLAGS) $< $(LOADLIBES) -o $@
```

You can redefine any of these macros by placing the appropriate macro definition at the beginning of the makefile.

You can redefine the action of a built-in rule by giving a new rule in your makefile. A built-in rule has the form

```
suffix-rule :
        command
```

where *suffix-rule* is a combination of suffixes showing the relationship of the implied target and dependent, and *command* is the XENIX command required

to carry out the rule. If more than one command is needed, they are given on separate lines.

The new rule must begin with an appropriate *suffix-rule*. The available *suffix-rules* are

.c	.c
.sh	.sh
.c.o	.c .o
.c .c	.s.o
.s .o	.y.o
.y .o	.l.o
.l .o	.y.c
.y .c	.l.c
.c.a	.c .a
.s .a	.h .h

A tilde ( ) indicates an SCCS file. A single suffix indicates a rule that makes an executable file from the given file. For example, the suffix rule ".c" is for the built-in rule that creates an executable file from a C source file. A pair of suffixes indicates a rule that makes one file from the other. For example, ".c.o" is for the rule that creates an object file (*.o*) file from a corresponding C source file (*.c*).

Any commands in the rule may use the built-in macros provided by **make**. For example, the following dependency line redefines the action of the *.c.o* rule.

```
.c.o :
        cc68 $< -c $*.o
```

If necessary, you can also create new *suffix-rules* by adding a list of new suffixes to a makefile with ".SUFFIXES". This pseudo-target name defines the suffixes that may be used to make *suffix-rules* for the built-in rules. The line has the form

.SUFFIXES: *suffix* ...

where *suffix* is usually a lowercase letter preceded by a dot (.). If more than one suffix is given, you must use spaces to separate them.

The order of the suffixes is significant. Each suffix is a dependent of the suffixes preceding it. For example, the suffix list

.SUFFIXES: .o .c .y .l .s

causes *prog.c* to be a dependent of *prog.o*, and *prog.y* to be a dependent of *prog.c*.

You can create new *suffix-rules* by combining dependent suffixes with the suffix of the intended target. The dependent suffix must appear first.

If a ".SUFFIXES" list appears more than once in a makefile, the suffixes are combined into a single list. If a ".SUFFIXES" is given that has no list, all suffixes are ignored.

4.9 Using Libraries

You can direct **make** to use a file contained in an archive library as a target or dependent. To do this you must explicitly name the file you wish to access by using a library name. A library name has the form

lib(*member-name*)

where *lib* is the name of the library containing the file, and *member-name* is the name of the file. For example, the library name

libtemp.a(print.o)

refers to the object file *print.o* in the archive library *libtemp.a*.

You can create your own built-in rules for archive libraries by adding the *.a* suffix to the suffix list, and creating new suffix combinations. For example, the combination ".c.a" may be used for a rule that defines how to create a library member from a C source file. Note that the dependent suffix in the new combination must be different than the suffix of the ultimate file. For example, the combination ".c.a" can be used for a rule that creates *.o* files, but not for one that creates *.c* files.

The most common use of the library naming convention is to create a makefile that automatically maintains an archive library. For example, the following dependency lines define the commands required to create a library, named *lib*, containing up to date versions of the files *file1.o*, *file2.o*, and *file3.o*.

```
lib:
        lib(file1.o) lib(file2.o) lib(file3.o)
        @echo lib is now up to date
.c.a:
        $(CC) -c $(CFLAGS) $<
        ar rv $@ $*.o
        rm -f $*.o
```

The *.c.a* rule shows how to redefine a built-in rule for a library. In the following example, the built-in rule is disabled, allowing the first dependency to create the library.

```
lib:
        lib(file1.o) lib(file2.o) lib(file3.o)
        $(CC) -c $(CFLAGS) $(?:.o=.c)
        ar rv lib $?
        rm $?
        @echo lib is now up to date
.c.a:;
```

In this example, a substitution sequence is used to change the value of the "$?" macro from the names of the object files "file1.o", "file2.o", and "file3.o" to "file1.c", "file2.c", and "file3.c".

4.10 Troubleshooting

Most difficulties in using **make** arise from **make**'s specific meaning of dependency. If the file *x.c* has the line

```
#include "defs"
```

then the object file *x.o* depends on *defs*; the source file *x.c* does not. (If *defs* is changed, it is not necessary to do anything to the file *x.c*, while it is necessary to recreate *x.o*.)

To determine which commands **make** will execute, without actually executing them, use the **–n** option. For example, the command

```
make –n
```

prints out the commands **make** would normally execute without actually executing them.

The debugging option **–d** causes **make** to print out a very detailed description of what it is doing, including the file times. The output is verbose, and recommended only as a last resort.

If a change to a file is absolutely certain to be benign (e.g., adding a new definition to an include file), the **–t** (touch) option can save a lot of time. Instead of issuing a large number of superfluous recompilations, **make** updates the modification times on the affected file. Thus, the command

```
make –ts
```

which stands for touch silently, causes the relevant files to appear up to date.

4.11 Using Make: An Example

As an example of the use of **make**, examine the makefile, given in Figure 4-1, used to maintain the **make** itself. The code for **make** is spread over a number

of C source files and a *yacc* grammar.

Make usually prints out each command before issuing it. The following output results from typing the simple command

```
make
```

in a directory containing only the source and makefile:

```
cc -c vers.c
cc -c main.c
cc -c doname.c
cc -c misc.c
cc -c files.c
cc -c dosys.c
yacc gram.y
mv y.tab.c gram.c
cc -c gram.c
cc vers.o main.o ... dosys.o gram.o -o make
13188+3348+3044 = 19580b = 046174b
```

Although none of the source files or grammars were mentioned by name in the makefile, **make** found them by using its suffix rules and issued the needed commands. The string of digits results from the **size make** command.

The last few targets in the makefile are useful maintenance sequences. The *print* target prints only the files that have been changed since the last **make print** command. A zero-length file, *print*, is maintained to keep track of the time of the printing; the **$?** macro in the command line then picks up only the names of the files changed since *print* was touched. The printed output can be sent to a different printer or to a file by changing the definition of the **P** macro.

Figure 4-1. Makefile Contents

```
# Description file for the make command

# Macro definitions below
P = lpr
FILES = Makefile vers.c defs main.c doname.c misc.c files.c dosys.c\
        gram.y lex.c
OBJECTS = vers.o main.o ... dosys.o gram.o
LIBES=
LINT = lint -p
CFLAGS = -O

#targets: dependents
#<TAB>actions

make:  $(OBJECTS)
        cc $(CFLAGS) $(OBJECTS) $(LIBES) -o make
        size make

$(OBJECTS):  defs
gram.o: lex.c

cleanup:
        -rm *.o gram.c
        -du

install:
        @size make /usr/bin/make
        cp make /usr/bin/make ; rm make

print:  $(FILES)  # print recently changed files
        pr $? | $P
        touch print

test:
        make -dp | grep -v TIME >1zap
        /usr/bin/make -dp | grep -v TIME >2zap
        diff 1zap 2zap
        rm 1zap 2zap

lint :  dosys.c doname.c files.c main.c misc.c vers.c gram.c
        $(LINT) dosys.c doname.c files.c main.c misc.c vers.c gram.c
        rm gram.c

arch:
        ar uv /sys/source/s2/make.a $(FILES)
```

Chapter 5
SCCS: A Source Code Control System

5.1 Introduction 5-1

5.2 Basic Information 5-1
- 5.2.1 Files and Directories 5-1
- 5.2.2 Deltas and SIDs 5-2
- 5.2.3 SCCS Working Files 5-3
- 5.2.4 SCCS Command Arguments 5-4
- 5.2.5 File Administrator 5-4

5.3 Creating and Using S-files 5-5
- 5.3.1 Creating an S-file 5-5
- 5.3.2 Retrieving a File for Reading 5-6
- 5.3.3 Retrieving a File for Editing 5-7
- 5.3.4 Saving a New Version of a File 5-8
- 5.3.5 Retrieving a Specific Version 5-9
- 5.3.6 Changing the Release Number of a File 5-9
- 5.3.7 Creating a Branch Version 5-10
- 5.3.8 Retrieving a Branch Version 5-10
- 5.3.9 Retrieving the Most Recent Version 5-11
- 5.3.10 Displaying a Version 5-11
- 5.3.11 Saving a Copy of a New Version 5-12
- 5.3.12 Displaying Helpful Information 5-12

5.4 Using Identification Keywords 5-13
- 5.4.1 Inserting a Keyword into a File 5-13
- 5.4.2 Assigning Values to Keywords 5-14
- 5.4.3 Forcing Keywords 5-14

5.5 Using S-file Flags 5-15
- 5.5.1 Setting S-file Flags 5-15
- 5.5.2 Using the i Flag 5-15
- 5.5.3 Using the d Flag 5-16

5.5.4 Using the v Flag 5-16
5.5.5 Removing an S-file Flag 5-16

5.6 Modifying S-file Information 5-16
5.6.1 Adding Comments 5-17
5.6.2 Changing Comments 5-17
5.6.3 Adding Modification Requests 5-18
5.6.4 Changing Modification Requests 5-18
5.6.5 Adding Descriptive Text 5-19

5.7 Printing from an S-file 5-20
5.7.1 Using a Data Specification 5-20
5.7.2 Printing a Specific Version 5-20
5.7.3 Printing Later and Earlier Versions 5-21

5.8 Editing by Several Users 5-21
5.8.1 Editing Different Versions 5-21
5.8.2 Editing a Single Version 5-22
5.8.3 Saving a Specific Version 5-22

5.9 Protecting S-files 5-23
5.9.1 Adding a User to the User List 5-23
5.9.2 Removing a User from a User List 5-23
5.9.3 Setting the Floor Flag 5-24
5.9.4 Setting the Ceiling Flag 5-24
5.9.5 Locking a Version 5-24

5.10 Repairing SCCS Files 5-25
5.10.1 Checking an S-file 5-25
5.10.2 Editing an S-file 5-25
5.10.3 Changing an S-file's Checksum 5-26
5.10.4 Regenerating a G-file for Editing 5-26
5.10.5 Restoring a Damaged P-file 5-26

5.11 Using Other Command Options 5-26
5.11.1 Getting Help With SCCS Commands 5-26
5.11.2 Creating a File With the Standard Input 5-27
5.11.3 Starting At a Specific Release 5-27
5.11.4 Adding a Comment to the First Version 5-27
5.11.5 Suppressing Normal Output 5-28
5.11.6 Including and Excluding Deltas 5-28

5.11.7 Listing the Deltas of a Version 5-29
5.11.8 Mapping Lines to Deltas 5-30
5.11.9 Naming Lines 5-30
5.11.10 Displaying a List of Differences 5-30
5.11.11 Displaying File Information 5-30
5.11.12 Removing a Delta 5-31
5.11.13 Searching for Strings 5-31
5.11.14 Comparing SCCS Files 5-32

5.1 Introduction

The Source Code Control System (SCCS) is a collection of XENIX commands that create, maintain, and control special files called SCCS files. The SCCS commands let you create and store multiple versions of a program or document in a single file, instead of one file for each version. The commands let you retrieve any version you wish at any time, make changes to this version, and save the changes as a new version of the file in the SCCS file.

The SCCS system is useful wherever you require a compact way to store multiple versions of the same file. The SCCS system provides an easy way to update any given version of a file and explicitly record the changes made. The commands are typically used to control changes to multiple versions of source programs, but may also be used to control multiple versions of manuals, specifications, and other documentation.

This chapter explains how to make SCCS files, how to update the files contained in SCCS files, and how to maintain the SCCS files once they are created. The following sections describe the basic information you need to start using the SCCS commands. Later sections describe the commands in detail.

5.2 Basic Information

This section provides some basic information about the SCCS system. In particular, it describes

- Files and directories
- Deltas and SIDs
- SCCS working files
- SCCS command arguments
- File administration

5.2.1 Files and Directories

All SCCS files (also called s-files) are originally created from text files containing documents or programs created by a user. The text files must have been created using a XENIX text editor such as **vi**. Special characters in the files are allowed only if they are also allowed by the given editor.

To simplify s-file storage, all logically related files (e.g., files belonging to the same project) should be kept in the same directory. Such directories should contain s-files only, and should have read and examine permission for everyone, and write permission for the user only.

Note that you must not use the XENIX **link** command to create multiple copies of an s-file.

5.2.2 Deltas and SIDs

Unlike an ordinary text file, an SCCS file (or s-file for short) contains nothing more than lists of changes. Each list corresponds to the changes needed to construct exactly one version of the file. The lists can then be combined to create the desired version from the original.

Each list of changes is called a "delta". Each delta has an identification string called an "SID". The SID is a string of at least two, and at most four, numbers separated by periods. The numbers name the version and define how it is related to other versions. For example, the first delta is usually numbered 1.1 and the second 1.2.

The first number in any SID is called the "release number". The release number usually indicates a group of versions that are similar and generally compatible. The second number in the SID is the "level number". It indicates major differences between files in the same release.

An SID may also have two optional numbers. The "branch number", the optional third number, indicates changes at a particular level, and the "sequence number", the fourth number, indicates changes at a particular branch. For example, the SIDs 1.1.1.1 and 1.1.1.2 indicate two new versions that contain slight changes to the original delta 1.1.

An s-file may at any time contain several different releases, levels, branches, and sequences of the same file. In general, the maximum number of releases an s-file may contain is 9999, that is, release numbers may range from 1 to 9999. The same limit applies to level, branch, and sequence numbers.

When you create a new version, the SCCS system usually creates a new SID by incrementing the level number of the original version. If you wish to create a new release, you must explicitly instruct the system to do so. A change to a release number indicates a major new version of the file. How to create a new version of a file and change release numbers is described later.

The SCCS system creates a branch and sequence number for the SID of a new version, if the next higher level number already exists. For example, if you change version 1.3 to create a version 1.4 and then change 1.3 again, the SCCS system creates a new version named 1.3.1.1.

Version numbers can become quite complicated. In general, it is wise to keep the numbers as simple as possible by carefully planning the creation of each new version.

5.2.3 SCCS Working Files

The SCCS system uses several different kinds of files to complete its tasks. In general, these files contain either actual text, or information about the commands in progress. For convenience, the SCCS system names these files by placing a prefix before the name of the original file from which all versions were made. The following is a list of the working files.

s-file — A permanent file that contains all versions of the given text file. The versions are stored as deltas, that is, lists of changes to be applied to the original file to create the given version. The name of an s-file is formed by placing the file prefix *s.* at the beginning of the original filename.

x-file — A temporary copy of the s-file. It is created by SCCS commands which change the s-file. It is used instead of the s-file to carry out the changes. When all changes are complete, the SCCS system removes the original s-file and gives the x-file the name of the original s-file. The name of the x-file is formed by placing the prefix *x.* at the beginning of the original file.

g-file — An ordinary text file created by applying the deltas in a given s-file to the original file. The g-file represents a copy of the given version of the original file, and as such receives the same filename as the original. When created, a g-file is placed in the current working directory of the user who requested the file.

p-file — A special file containing information about the versions of an s-file currently being edited. The p-file is created when a g-file is retrieved from the s-file. The p-file exists until all currently retrieved files have been saved in the s-file; it is then deleted. The p-file contains one or more entries describing the SID of the retrieved g-file, the proposed SID of the new, edited g-file, and the login name of the user who retrieved the g-file. The p-file name is formed by placing the prefix *p.* at the beginning of the original filename.

z-file — A lock file used by SCCS commands to prevent two users from updating a single SCCS file at the same time. Before a command modifes an SCCS file, it creates a z-file and copies its own process ID to it. Any other command which attempts to access the file while the z-file is present displays an error message and stops. When the original command has finished its tasks, it deletes the z-file before stopping. The z-file name is formed by placing the prefix *z.* at the beginning of the original filename.

l-file — A special file containing a list of the deltas required to create a given version of a file. The l-file name is formed by placing the prefix *l.* at the beginning of the original filename.

d-file — A temporary copy of the g-file used to generate a new delta.

q-file — A temporary file used by the **delta** command when updating the p-file. The file is not directly accessible.

In general, a user never directly accesses x-files, z-files, d-files, or q-files. If a system crash or similar situation abnormally terminates a command, the user may wish delete these files to ensure proper operation of subsequent SCCS commands.

5.2.4 SCCS Command Arguments

Almost all SCCS commands accept two types of arguments: options and filenames. These appear in the SCCS command line immediately after the command name.

An option indicates a special action to be taken by the given SCCS command. An option is usually a lowercase letter preceded by a minus sign (–). Some options require an additional name or value.

A filename indicates the file to be acted on. The syntax for SCCS filenames is like other XENIX filename syntax. Appropriate pathnames must be given if required. Some commands also allow directory names. In this case, all files in the directory are acted on. If the directory contains non-SCCS and unreadable files, these are ignored. A filename must not begin with a minus sign (–).

The special symbol – may be used to cause the given command to read a list of filenames from the standard input. These filenames are then used as names for the files to be processed. The list must terminate with an end-of-file character.

Any options given with a command apply to all files. The SCCS commands process the options before any filenames, so the options may appear anywhere on the command line.

Filenames are processed left to right. If a command encounters a fatal error, it stops processing the current file and, if any other files have been given, begins processing the next.

5.2.5 File Administrator

Every SCCS file requires an administrator to maintain and keep the file in order. The administrator is usually the user who created the file and therefore owns it. Before other users can access the file, the administrator must ensure that they have adequate access. Several SCCS commands let the administrator define who has access to the versions in a given s-file. These are described later.

5.3 Creating and Using S-files

The s-file is the key element in the SCCS system. It provides compact storage for all versions of a given file and automatic maintenance of the relationships between the versions.

This section explains how to use the **admin**, **get**, and **delta** commands to create and use s-files. In particular, it describes how to create the first version of a file, how to retrieve versions for reading and editing, and how to save new versions.

5.3.1 Creating an S-file

You can create an s-file from an existing text file using the **–i** (for "initialize") option of the **admin** command. The command has the form

admin –i*filename* *s.filename*

where –i*filename* gives the name of the text file from which the s-file is to be created, and *s.filename* is the name of the new s-file. The name must begin with *s.* and must be unique; no other s-file in the same directory may have the same name. For example, suppose the file named *demo.c* contains the short C language program

```
#include <stdio.h>

main ()
{
printf("This is version 1.1 \n");
}
```

To create an s-file, type

```
admin –idemo.c   s.demo.c
```

This command creates the s-file *s.demo.c*, and copies the first delta describing the contents of *demo.c* to this new file. The first delta is numbered **1.1**.

After creating an s-file, the original text file should be removed using the **rm** command, since it is no longer needed. If you wish to view the text file or make changes to it, you can retrieve the file using the **get** command described in the next section.

When first creating an s-file, the **admin** command may display the warning message

```
No id keywords (cm7)
```

In general, this message can be ignored unless you have specifically included keywords in your file (see the section, "Using Identification Keywords" later in this chapter).

Note that only a user with write permission in the directory containing the s-file may use the **admin** command on that file. This protects the file from administration by unauthorized users.

5.3.2 Retrieving a File for Reading

You can retrieve a file for reading from a given s-file by using the **get** command. The command has the form

get *s.filename* ...

where *s.filename* is the name of the s-file containing the text file. The command retrieves the lastest version of the text file and copies it to a regular file. The file has the same name as the s-file but with the *s.* removed. It also has read-only file permissions. For example, suppose the s-file *s.demo.c* contains the first version of the short C program shown in the previous section. To retrieve this program, type

```
get s.demo.c
```

The command retrieves the program and copies it to the file named *demo.c*. You may then display the file just as you do any other text file.

The command also displays a message which describes the SID of the retrieved file and its size in lines. For example, after retrieving the short C program from *s.demo.c*, the command displays the message

```
1.1
6 lines
```

You may also retrieve more than one file at a time by giving multiple s-file names in the command line. For example, the command

```
get s.demo.c s.def.h
```

retrieves the contents of the s-files *s.demo.c* and *s.def.h* and copies them to the text files *demo.c* and *def.h*. When giving multiple s-file names in a command, you must separate each with at least one space. When the **get** command displays information about the files, it places the corresponding filename before the relevent information.

5.3.3 Retrieving a File for Editing

You can retrieve a file for editing from a given s-file by using the –e (for "editing") option of the **get** command. The command has the form

get –e *s.filename* ...

where *s.filename* is the name of the s-file containing the text file. You may give more than one filename if you wish. If you do, you must separate each name with a space.

The command retrieves the lastest version of the text file and copies it to an ordinary text file. The file has the same name as the s-file but with the *s.* removed. It has read and write file permissions. For example, suppose the s-file *s.demo.c* contains the first version of a C program. To retrieve this program, type

get –e s.demo.c

The command retrieves the program and copies it to the file named *demo.c*. You may edit the file just as you do any other text file.

If you give more than one filename, the command creates files for each corresponding s-file. Since the –e option applies to all the files, you may edit each one.

After retrieving a text file, the command displays a message giving the SID of the file and its size in lines. The message also displays a proposed SID, that is, the SID for the new version after editing. For example, after retrieving the six-line C program in *s.demo.c*, the command displays the message

1.1
new delta 1.2
6 lines

The proposed SID is 1.2. If more than one file is retrieved, the corresponding filename precedes the relevant information.

Note that any changes made to the text file are not immediately copied to the corresponding s-file. To save these changes you must use the **delta** command described in the next section. To help keep track of the current file version, the **get** command creates another file, called a p-file, that contains information about the text file. This file is used by a subsequent **delta** command when saving the new version. The p-file has the same name as the s-file but begins with a *p.*. The user must not access the p-file directly.

5.3.4 Saving a New Version of a File

You can save a new version of a text file by using the **delta** command. The command has the form

delta *s.filename*

where *s.filename* is the name of the s-file from which the modified text file was retrieved. For example, to save changes made to a C program in the file *demo.c* (which was retrieved from the file *s.demo.c*), type

```
delta s.demo.c
```

Before saving the new version, the **delta** command asks for comments explaining the nature of the changes. It displays the prompt

```
comments?
```

You may type any text you think appropriate, up to 512 characters. The comment must end with a newline character. If necessary, you can start a new line by typing a backslash (\) followed by a newline character. If you do not wish to include a comment, just type a newline character.

Once you have given a comment, the command uses the information in the corresponding p-file to compare the original version with the new version. A list of all the changes is copied to the s-file. This is the new delta.

After a command has copied the new delta to the s-file, it displays a message showing the new SID and the number of lines inserted, deleted, or left unchanged in the new version. For example, if the C program has been changed to

```
#include <stdio.h>

main ()
{
int i = 2;

printf("This is version 1.%d 0, i);
}
```

the command displays the message

```
1.2
3 inserted
1 deleted
5 unchanged
```

Once a new version is saved, the next **get** command retrieves the new version.

The command ignores previous versions. If you wish to retrieve a previous version, you must use the –r option of the **get** command as described in the next section.

5.3.5 Retrieving a Specific Version

You can retrieve any version you wish from an s-file by using the –r (for "retrieve") of the **get** command. The command has the form

get [–e] –r*SID* *s.filename* ...

where –e is the edit option, –r*SID* gives the SID of the version to be retrieved, and *s.filename* is the name of the s-file containing the file to be retrieved. You may give more than one filename. The names must be separated with spaces.

The command retrieves the given version and copies it to the file having the same name as s-file but with the *s.* removed. The file has read-only permission unless you also give the –e option. If multiple filenames are given, one text file of the given version is retrieved from each. For example, the command

```
get –r1.1 s.demo.c
```

retrieves version 1.1 from the s-file *s.demo.c*, but the command

```
get –e –r1.1 s.demo.c  s.def.h
```

retrieves for editing a version 1.1 from both *s.demo.c* and *s.def.h*. If you give the number of a version that does not exist, the command displays an error message.

You may omit the level number of a version number if you wish, that is, just give a release number. If you do, the command automatically retrieves the most recent version having the same release number. For example, if the most recent version in the file *s.demo.c* is numbered 1.4, the command

```
get –r1 s.demo.c
```

retrieves the version 1.4. If there is no version with the given release number, the command retrieves the most recent version in the previous release.

5.3.6 Changing the Release Number of a File

You can direct the **delta** command to change the release number of a new version of a file by using the –r option of the **get** command. In this case, the **get** command has the form

get –e –r*rel-num* *s.filename* ...

where –e is the required edit option, –r*rel-num* gives the new release number of the file, and *s.filename* gives the name of the s-file containing the file to be retrieved. The new release number must be an entirely new number, that is, no existing version may have this number. You may give more than one filename.

The command retrieves the most recent version from the s-file, then copies the new release number to the p-file. On the subsequent **delta** command, the new version is saved using the new release number and level number 1. For example, if the most recent version in the s-file *s.demo.c* is 1.4, the command

```
get –e –r2 s.demo.c
```

causes the subsequent **delta** to save a new version 2.1, not 1.5. The new release number applies to the new version only; the release numbers of previous versions are not affected. Therefore, if you edit version 1.4 (from which 2.1 was derived) and save the changes, you create a new version 1.5. Similarly, if you edit version 2.1, you create a new version 2.2.

As before, the **get** command also displays a message showing the current version number, the proposed version number, and the size of the file in lines. Similarly, the subsequent **delta** command displays the new version number and the number of lines inserted, deleted, and unchanged in the new file.

5.3.7 Creating a Branch Version

You can create a branch version of a file by editing a version that has been previously edited. A branch version is simply a version whose SID contains a branch and sequence number.

For example, if version 1.4 already exists, the command

```
get –e –r1.3 s.demo.c
```

retrieves version 1.3 for editing and gives 1.3.1.1 as the proposed SID.

In general, whenever **get** discovers that you wish to edit a version that already has a succeeding version, it uses the first available branch and sequence numbers for the proposed SID. For example, if you edit version 1.3 a third time, **get** gives 1.3.2.1 as the proposed SID.

You can save a branch version just like any other version by using the **delta** command.

5.3.8 Retrieving a Branch Version

You can retrieve a branch version of a file by using the –r option of the **get** command. For example, the command

```
get -r1.3.1.1 s.demo.c
```

retrieves branch version 1.3.1.1.

You may retrieve a branch version for editing by using the **-e** option of the **get** command. When retrieving for editing, **get** creates the proposed SID by incrementing the sequence number by one. For example, if you retrieve branch version 1.3.1.1 for editing, **get** gives 1.3.1.2 as the proposed SID.

As always, the command displays the version number and file size. If the given branch version does not exist, the command displays an error message.

You may omit the sequence number if you wish. In this case, the command retrieves the most recent branch version with the given branch number. For example, if the most recent branch version in the s-file *s.def.h* is 1.3.1.4, the command

```
get -r1.3.1 s.def.h
```

retrieves version 1.3.1.4.

5.3.9 Retrieving the Most Recent Version

You can always retrieve the most recent version of a file by using the **-t** option with the **get** command. For example, the command

```
get -t s.demo.c
```

retrieves the most recent version from the file *s.demo.c*. You may combine the **-r** and **-t** options to retrieve the most recent version of a given release number. For example, if the most recent version with release number 3 is 3.5, then the command

```
get -r3 -t s.demo.c
```

retrieves version 3.5. If a branch version exists that is more recent than version 3.5 (e.g., 3.2.1.5), then the above command retrieves the branch version and ignores version 3.5.

5.3.10 Displaying a Version

You can display the contents of a version at the standard output by using the **-p** option of the **get** command. For example, the command

```
get -p s.demo.c
```

displays the most recent version in the s-file *s.demo.c* at the standard output. Similarly, the command

```
get -p -r2.1 s.demo.c
```

displays version 2.1 at the standard output.

The -p option is useful for creating g-files with user-supplied names. This option also directs all output normally sent to the standard output, such as the SID of the retrieved file, to the standard error file. Thus, the resulting file contains only the contents of the given version. For example, the command

```
get -p s.demo.c >version.c
```

copies the most recent version in the s-file *s.demo.c* to the file *version.c*. The SID of the file and its size is copied to the standard error file.

5.3.11 Saving a Copy of a New Version

The **delta** command normally removes the edited file after saving it in the s-file. You can save a copy of this file by using the -n option of the **delta** command. For example, the command

```
delta -n s.demo.c
```

first saves a new version in the s-file *s.demo.c*, then saves a copy of this version in the file *demo.c*. You may display the file as desired, but you cannot edit the file.

5.3.12 Displaying Helpful Information

An SCCS command displays an error message whenever it encounters an error in a file. An error message has the form

ERROR [*filename*]: *message* (*code*)

where *filename* is the name of the file being processed, *message* is a short description of the error, and *code* is the error code.

You may use the error code as an argument to the **help** command to display additional information about the error. The command has the form

```
help code
```

where *code* is the error code given in an error message. The command displays one or more lines of text that explain the error and suggest a possible remedy. For example, the command

```
help co1
```

displays the message

```
co1:
"not an SCCS file"
A file that you think is an SCCS file
does not begin with the characters "s.".
```

The **help** command can be used at any time.

5.4 Using Identification Keywords

The SCCS system provides several special symbols, called identification keywords, which may be used in the text of a program or document to represent a predefined value. Keywords represent a wide range of values, from the creation date and time of a given file, to the name of the module containing the keyword. When a user retrieves the file for reading, the SCCS system automatically replaces any keywords it finds in a given version of a file with the keyword's value.

This section explains how keywords are treated by the various SCCS commands, and how you may use the keywords in your own files. Only a few keywords are described in this section. For a complete list of the keywords, see the section *get*(CP) in the XENIX *Reference Manual*.

5.4.1 Inserting a Keyword into a File

You may insert a keyword into any text file. A keyword is simply an uppercase letter enclosed in percent signs (%). No special characters are required. For example, "%I%" is the keyword representing the SID of the current version, and "%H%" is the keyword representing the current date.

When the program is retrieved for reading using the **get** command, the keywords are replaced by their current values. For example, if the "%M%", "%I%", and "%H" keywords are used in place of the module name, the SID, and the current data in a program statement

```
char header(100) = {" %M%  %I%  %H% "};
```

then these keywords are expanded in the retrieved version of the program

```
char header(100) = {" MODNAME  2.3  07/07/77 "};
```

The **get** command does not replace keywords when retrieving a version for editing. The system assumes that you wish keep the keywords (and not their values) when you save the new version of the file.

To indicate that a file has no keywords, the **get**, **delta**, and **admin** commands display the message

No id keywords (cm7)

This message is normally treated as a warning, letting you know that no keywords are present. However, you may change the operation of the system to make this a fatal error, as explained later in this chapter.

5.4.2 Assigning Values to Keywords

The values of most keywords are predefined by the system, but some, such as the value for the "%M%" keyword can be explicitly defined by the user. To assign a value to a keyword, you must set the corresponding s-file flag to the desired value. You can do this by using the **–f** option of the **admin** command.

For example, to set the %M% keyword to "cdemo", you must set the **m** flag as in the command

admin –fmcdemo s.demo.c

This command records "cdemo" as the current value of the %M% keyword. Note that if you do not set the **m** flag, the SCCS system uses the name of the original text file for %M% by default.

The **t** and **q** flags are also associated with keywords. A description of these flags and the corresponding keywords can be found in the section *get*(CP) in the XENIX *Reference Manual*. You can change keyword values at any time.

5.4.3 Forcing Keywords

If a version is found to contain no keywords, you can force a fatal error by setting the **i** flag in the given s-file. The flag causes the **delta** and **admin** commands to stop processing of the given version and report an error. The flag is useful for ensuring that keywords are used properly in a given file.

To set the **i** flag, you must use the **–f** option of the **admin** command. For example, the command

admin –fi s.demo.c

sets the **i** flag in the s-file *s.demo.c*. If the given version does not contain keywords, subsequent **delta** or **admin** commands that access this file print an error message.

Note that if you attempt to set the **i** flag at the same time as you create an s-file, and if the initial text file contains no keywords, the **admin** command displays a fatal error message and stops without creating the s-file.

5.5 Using S-file Flags

An s-file flag is a special value that defines how a given SCCS command will operate on the corresponding s-file. The s-file flags are stored in the s-file and are read by each SCCS command before it operates on the file. S-file flags affect operations such as keyword checking, keyword replacement values, and default values for commands.

This section explains how to set and use s-file flags. It also describes the action of commonly-used flags. For a complete description of all flags, see the section *admin*(CP) in the XENIX *Reference Manual*.

5.5.1 Setting S-file Flags

You can set the flags in a given s-file by using the **–f** option of the **admin** command. The command has the form

admin –f*flag* *s.filename*

where –f*flag* gives the flag to be set, and *s.filename* gives the name of the s-file in which the flag is to be set. For example, the command

admin –fi s.demo.c

sets the **i** flag in the s-file *s.demo.c*.

Note that some s-file flags take values when they are set. For example, the **m** flag requires that a module name be given. When a value is required, it must immediately follow the flag name, as in the command

admin –fmdmod s.demo.c

which sets the **m** flag to the module name "dmod".

5.5.2 Using the i Flag

The **i** flag causes the **admin** and **delta** commands to print a fatal error message and stop, if no keywords are found in the given text file. The flag is used to prevent a version of a file, which contains expanded keywords, from being saved as a new version. (Saving an expanded version destroys the keywords for all subsequent versions).

When the **i** flag is set, each new version of a file must contain at least one keyword. Otherwise, the version cannot be saved.

5.5.3 Using the d Flag

The **d** flag gives the default SID for versions retrieved by the **get** command. The flag takes an SID as its value. For example, the command

```
admin -fd1.1 s.demo.c
```

sets the default SID to 1.1. A subsequent **get** command which does not use the **-r** option will retrieve version 1.1.

5.5.4 Using the v Flag

The **v** flag allows you to include modification requests in an s-file. Modification requests are names or numbers that may be used as a shorthand means of indicating the reason for each new version.

When the **v** flag is set, the **delta** command asks for the modification requests just before asking for comments. The **v** flag also allows the **-m** option to be used in the **delta** and **admin** commands.

5.5.5 Removing an S-file Flag

You can remove an s-file flag from an s-file by using the **-d** option of the **admin** command. The command has the form

```
admin -dflag  s.filename
```

where **-d***flag* gives the name of the flag to be removed and *s.filename* is the name of the s-file from which the flag is to be removed. For example, the command

```
admin -di s.demo.c
```

removes the **i** flag from the s-file *s.demo.c*. When removing a flag which takes a value, only the flag name is required. For example, the command

```
admin -dm s.demo.c
```

removes the **m** flag from the s-file.

The **-d** and **-i** options must not be used at the same time.

5.6 Modifying S-file Information

Every s-file contains information about the deltas it contains. Normally, this information is maintained by the SCCS commands and is not directly accessible

by the user. Some information, however, is specific to the user who creates the s-file, and may be changed as desired to meet the user's requirements. This information is kept in two special parts of the s-file called the "delta table" and the "description field".

The delta table contains information about each delta, such as the SID and the date and time of creation. It also contains user-supplied information, such as comments and modification requests. The description field contains a user-supplied description of the s-file and its contents. Both parts can be changed or deleted at any time to reflect changes to the s-file contents.

5.6.1 Adding Comments

You can add comments to an s-file by using the –y option of the **delta** and **admin** commands. This option causes the given text to be copied to the s-file as the comment for the new version. The comment may be any combination of letters, digits, and punctuation symbols. No embedded newline characters are allowed. If spaces are used, the comment must be enclosed in double quotes. The complete command must fit on one line. For example, the command

```
delta –y"George Wheeler"  s.demo.c
```

saves the comment "George Wheeler" in the s-file *s.demo.c*.

The –y option is typically used in shell procedures as part of an automated approach to maintaining files. When the option is used, the **delta** command does not print the corresponding comment prompt, so no interaction is required. If more than one s-file is given in the command line, the given comment applies to them all.

5.6.2 Changing Comments

You can change the comments in a given s-file by using the **cdc** command. The command has the form

```
cdc –rSID  s.filename
```

where –r*SID* gives the SID of the version whose comment is to be changed, and *s.filename* is the name of the s-file containing the version. The command asks for a new comment by displaying the prompt

```
comments?
```

You may type any sequence of characters up to 512 characters long. The sequence may contain embedded newline characters if they are preceded by a backslash (\). The sequence must be terminated with a newline character. For example, the command

```
cdc -r3.4 s.demo.c
```

prompts for a new comment for version 3.4.

Although the command does not delete the old comment, it is no longer directly accessible by the user. The new comment contains the login name of the user who invoked the **cdc** command and the time the comment was changed.

5.6.3 Adding Modification Requests

You can add modification requests to an s-file, when the **v** flag is set, by using the **-m** option of the **delta** and **admin** commands. A modification request is a shorthand method of describing the reason for a particular version. Modification requests are usually names or numbers which the user has chosen to represent a specific request.

The **-m** option causes the given command to save the requests following the option. A request may be any combination of letters, digits, and punctuation symbols. If you give more than one request, you must separate them with spaces and enclose the request in double quotes. For example, the command

```
delta -m"error35 optimize10" s.demo.c
```

copies the requests "error35" and "optimize10" to *s.demo.c*, while saving the new version.

The **-m** option, when used with the **admin** command, must be combined with the **-i** option. Furthermore, the **v** flag must be explicitly set with the **-f** option. For example, the command

```
admin -idef.h -m"error0" -fv s.def.h
```

inserts the modification request "error0" in the new file *s.def.h*.

The **delta** command does not prompt for modification requests if you use the **-m** option.

5.6.4 Changing Modification Requests

You can change modification requests, when the **v** flag is set, by using the **cdc** command. The command asks for a list of modification requests by displaying the prompt

```
MRs?
```

You may type any number of requests. Each request may have any combination of letters, digits, or punctuation symbols. No more than 512 characters are allowed, and the last request must be terminated with a newline

character. If you wish to remove a request, you must precede the request with an exclamation mark (!). For example, the command

```
cdc -r1.4 s.demo.c
```

asks for changes to the modification requests. The response

```
MRs? error36 !error35
```

adds the request "error36" and removes "error35".

5.6.5 Adding Descriptive Text

You can add descriptive text to an s-file by using the –t option of the **admin** command. Descriptive text is any text that describes the purpose and reason for the given s-file. Descriptive text is independent of the contents of the s-file and can only be displayed using the **prs** command.

The –t option directs the **admin** to copy the contents of a given file into the description field of the s-file. The command has the form

```
admin -tfilename  s.filename
```

where –t*filename* gives the name of the file containing the descriptive text, and *s.filename* is the name of the s-file to receive the descriptive text. The file to be inserted may contain any amount of text. For example, the command

```
admin -tcdemo s.demo.c
```

inserts the contents of the file *cdemo* into the description field of the s-file *s.demo.c*.

The –t option may also be used to initialize the description field when creating the s-file. For example, the command

```
admin -idemo.c -tcdemo s.demo.c
```

inserts the contents of the file *cdemo* into the new s-file *s.demo.c*. If –t is not used, the description field of the new s-file is left empty.

You can remove the current descriptive text in an s-file by using the –t option without a filename. For example, the command

```
admin  -t  s.demo.c
```

removes the descriptive text from the s-file *s.demo.c*.

5.7 Printing from an S-file

This section explains how to use the **prs** command to display information contained in an s-file. The **prs** command has a variety of options which control the display format and content.

5.7.1 Using a Data Specification

You can explicitly define the information to be printed from an s-file by using the **–d** option of the **prs** command. The command copies user-specified information to the standard output. The command has the form

prs –d*spec* *s.filename*

where –d*spec* is the data specification, and *s.filename* is the name of the s-file from which the information is to be taken.

The data specification is a string of data keywords and text. A data keyword is an uppercase letter, enclosed in colons (:). It represents a value contained in the given s-file. For example, the keyword **:I:** represents the SID of a given version, **:F:** represent the filename of the given s-file, **:C:** represents the comment line associated with a given version. Data keywords are replaced by these values when the information is printed.

For example, the command

prs –d" version: :I: filename: :F: " s.demo.c

may produce the line

version: 2.1 filename: s.demo.c

A complete list of the data keywords is given in the section *prs*(CP) in the XENIX *Reference Manual*.

5.7.2 Printing a Specific Version

You can print information about a specific version in a given s-file by using the **–r** option of the **prs** command. The command has the form

prs –r*SID* *s.filename*

where –r*SID* gives the SID of the desired version, and *s.filename* is the name of the s-file containing the version. For example, the command

prs –r2.1 s.demo.c

prints information about version 2.1 in the s-file *s.demo.c*.

If the –r option is not specified, the command prints information about the most recently created delta.

5.7.3 Printing Later and Earlier Versions

You can print information about a group of versions by using the –l and –e options of the **prs** command. The –l option causes the command to print information about all versions immediately succeeding the given version. The –e option causes the command to print information about all versions immediately preceding the given version. For example, the command

```
prs –r1.4 –e s.demo.c
```

prints all information about versions which precede version 1.4 (e.g., 1.3, 1.2, and 1.1). The command

```
prs –r1.4 –l s.abc
```

prints information about versions which succeed version 1.4 (e.g., 1.5, 1.6, and 2.1).

If both options are given, information about all versions is printed.

5.8 Editing by Several Users

The SCCS system allows any number users to access and edit versions of a given s-file. Since users are likely to access different versions of the s-file at the same time, the system is designed to allow concurrent editing of different versions. Normally, the system allows only one user at a time to edit a given version, but you can allow concurrent editing of the same version by setting the **j** flag in the given s-file.

The following sections explain how to perform concurrent editing and how to save edited versions when you have retrieved more than one version for editing.

5.8.1 Editing Different Versions

The SCCS system allows several different versions of a file to be edited at the same time. This means a user can edit version 2.1 while another user edit version 1.1. There is no limit to the number of versions which may be edited at any given time.

When several users edits different versions concurrently, each user must begin work in his own directory. If users attempt to share a directory and work on versions from the same s-file at the same time, the **get** command will refuse to

retrieve a version.

5.8.2 Editing a Single Version

You can let a single version of a file be edited by more than one user by setting the **j** flag in the given s-file. The flag causes the **get** command to check the p-file and create a new proposed SID if the given version is already being edited.

You can set the flag by using the **–f** option of the **admin** command. For example, the command

```
admin –fj s.demo.c
```

sets the flag for the s-file *s.demo.c*.

When the flag is set, the **get** command uses the next available branch SID for each new proposed SID. For example, suppose a user retrieves for editing version 1.4 in the file *s.demo.c*, and that the proposed version is 1.5. If another user retrieves version 1.4 for editing before the first user has saved his changes, the the proposed version for the new user will be 1.4.1.1, since version 1.5 is already proposed and likely to be taken. In no case will a version edited by two separate users result in a single new version.

5.8.3 Saving a Specific Version

When editing two or more versions of a file, you can direct the **delta** command to save a specific version by using the **–r** option to give the SID of that version. The command has the form

```
delta –rSID  s.filename
```

where **–r***SID* gives the SID of the version being saved, and *s.filename* is the name of the s-file to receive the new version. The *SID* may be the SID of the version you have just edited, or the proposed SID for the new version. For example, if you have retrieved version 1.4 for editing (and no version 1.5 exists), both commands

```
delta –r1.5 s.demo.c
```

and

```
delta –r1.4 s.demo.c
```

save version 1.5.

5.9 Protecting S-files

The SCCS system uses the normal XENIX system file permissions to protect s-files from changes by unauthorized users. In addition to the XENIX system protections, the SCCS system provides two ways to protect the s-files: the "user list" and the "protection flags". The user list is a list of login names and group IDs of users who are allowed to access the s-file and create new versions of the file. The protection flags are three special s-file flags that define which versions are currently accessible to otherwise authorized users. The following sections explain how to set and use the user list and protection flags.

5.9.1 Adding a User to the User List

You can add a user or a group of users to the user list of a given s-file by using the **–a** option of the **admin** command. The option causes the given name to be added to the user list. The user list defines who may access and edit the versions in the s-file. The command has the form

admin –a*name* *s.filename*

where –a*name* gives the login name of the user or the group name of a group of users to be added to the list, and *s.filename* gives the name of the s-file to receive the new users. For example, the command

admin –ajohnd –asuex –amarketing s.demo.c

adds the users "johnd" and "suex" and the group "marketing" to the user list of the s-file *s.demo.c*.

If you create an s-file without giving the –a option, the user list is left empty, and all users may access and edit the files. When you explicitly give a user name or names, only those users can access the files.

5.9.2 Removing a User from a User List

You can remove a user or a group of users from the user list of a given s-file by using the –e option of the **admin** command. The option is similar to the **–a** option but performs the opposite operation. The command has the form

admin –e*name* *s.filename*

where –e*name* gives the login name of a user or the group name of a group of users to be removed from the list, and *s.filename* is the name of the s-file from which the names are to be removed. For example, the command

admin –ejohnd –emarketing s.demo.c

removes the user "johnd" and the group "marketing" from the user list of the s-file *s.demo.c*.

5.9.3 Setting the Floor Flag

The floor flag, **f**, defines the release number of the lowest version a user may edit in a given s-file. You can set the flag by using the **–f** option of the **admin** command. For example, the command

```
admin –ff2 s.demo.c
```

sets the floor to release number 2. If you attempt to retrieve any versions with a release number less than 2, an error will result.

5.9.4 Setting the Ceiling Flag

The ceiling flag, **c**, defines the release number of the highest version a user may edit in a given s-file. You can set the flag by using the **–f** option of the **admin** command. For example, the command

```
admin –fc5 s.demo.c
```

sets the ceiling to release number 5. If you attempt to retrieve any versions with a release number greater than 5, an error will result.

5.9.5 Locking a Version

The lock flag, **l**, lists by release number all versions in a given s-file which are locked against further editing. You can set the flag by using the **–f** flag of the **admin** command. The flag must be followed by one or more release numbers. Multiple release numbers must be separated by commas (,). For example, the command

```
admin –fl3 s.demo.c
```

locks all versions with release number 3 against further editing. The command

```
admin –fl4,5,9 s.def.h
```

locks all versions with release numbers 4, 5, and 9.

Note that the special symbol "a" may be used to specify all release numbers. The command

```
admin –fla  s.demo.c
```

locks all versions in the file *s.demo.c*.

5.10 Repairing SCCS Files

The SCCS system carefully maintains all SCCS files, making damage to the files very rare. However, damage can result from hardware malfunctions, which cause incorrect information to be copied to the file. The following sections explain how to check for damage to SCCS files, and how to repair the damage or regenerate the file.

5.10.1 Checking an S-file

You can check a file for damage by using the **–h** option of the **admin** command. This option causes the checksum of the given s-file to be computed and compared with the existing sum. An s-file's checksum is an internal value computed from the sum of all bytes in the file. If the new and existing checksums are not equal, the command displays the message

```
corrupted file (co6)
```

indicating damage to the file. For example, the command

```
admin –h s.demo.c
```

checks the s-file *s.demo.c* for damage by generating a new checksum for the file, and comparing the new sum with the existing sum.

You may give more than one filename. If you do, the command checks each file in turn. You may also give the name of a directory, in which case, the command checks all files in the directory.

Since failure to repair a damaged s-file can destroy the file's contents or make the file inaccessible, it is a good idea to regularly check all s-files for damage.

5.10.2 Editing an S-file

When an s-file is discovered to be damaged, it is a good idea to restore a backup copy of the file from a backup disk rather than attempting to repair the file. (Restoring a backup copy of a file is described in the XENIX *Operations Guide*.) If this is not possible, the file may be edited using a XENIX text editor.

To repair a damaged s-file, use the description of an s-file given in the section *sccsfile*(F) in the XENIX *Reference Manual*, to locate the part of the file which is damaged. Use extreme care when making changes; small errors can cause unwanted results.

5.10.3 Changing an S-file's Checksum

After repairing a damaged s-file, you must change the file's checksum by using the –z option of the **admin** command. For example, to restore the checksum of the repaired file *s.demo.c*, type

```
admin –z s.demo.c
```

The command computes and saves the new checksum, replacing the old sum.

5.10.4 Regenerating a G-file for Editing

You can create a g-file for editing without affecting the current contents of the p-file by using the –**k** option of the **get** command. The option has the same affect as the –e option, except that the current contents of the p-file remain unchanged. The option is typically used to regenerate a g-file that has been accidentally removed or destroyed before it has been saved using the **delta** command.

5.10.5 Restoring a Damaged P-file

The –**g** option of the **get** command may be used to generate a new copy of a p-file that has been accidentally removed. For example, the command

```
get –e –g s.demo.c
```

creates a new p-file entry for the most recent version in *s.demo.c*. If the file *demo.c* already exists, it will not be changed by this command.

5.11 Using Other Command Options

Many of the SCCS commands provide options that control their operation in useful ways. This section describes these options and explains how you may use them to perform useful work.

5.11.1 Getting Help With SCCS Commands

You can display helpful information about an SCCS command by giving the name of the command as an argument to the **help** command. The **help** command displays a short explanation of the command and command syntax. For example, the command

```
help rmdel
```

displays the message

rmdel:

```
rmdel –rSID name ...
```

5.11.2 Creating a File With the Standard Input

You can direct **admin** to use the standard input as the source for a new s-file by using the –i option without a filename. For example, the command

```
admin –i s.demo.c <demo.c
```

causes **admin** to create a new s-file named *s.demo.c* which uses the text file *demo.c* as its first version.

This method of creating a new s-file is typically used to connect **admin** to a pipe. For example, the command

```
cat mod1.c mod2.c | admin –i s.mod.c
```

creates a new s-file *s.mod.c* which contains the first version of the concatenated files *mod1.c* and *mod2.c*.

5.11.3 Starting At a Specific Release

The **admin** command normally starts numbering versions with release number 1. You can direct the command to start with any given release number by using the –r option. The command has the form

```
admin –rrel-num s.filename
```

where –r*rel-num* gives the value of the starting release number, and *s.filename* is the name of the s-file to be created. For example, the command

```
admin –idemo.c –r3 s.demo.c
```

starts with release number 3. The first version is 3.1.

5.11.4 Adding a Comment to the First Version

You can add a comment to the first version of file by using the –y option of the **admin** command when creating the s-file. For example, the command

```
admin –idemo.c –y"George Wheeler" s.demo.c
```

inserts the comment "George Wheeler" in the new s-file *s.demo.c*.

The comment may be any combination of letters, digits, and punctuation symbols. If spaces are used, the comment must be enclosed in double quotes. The complete command must fit on one line.

If the –y option is not used when creating an s-file, a comment of the form

date and time created YY/MM/DD HH:MM:SS by logname

is automatically inserted.

5.11.5 Suppressing Normal Output

You can suppress the normal display of messages created by the **get** command by using the –s option. The option prevents information, such as the SID of the retrieved file, from being copied to the standard output. The option does not suppress error messages.

The –s option is often used with the –p option to pipe the output of the **get** command to other commands. For example, the command

```
get –p –s s.demo.c | lpr
```

copies the most recent version in the s-file *s.demo.c* to the line printer.

You can also suppress the normal output of the **delta** command by using the –s option. This option suppresses all output normally directed to the standard output, except for the normal comment prompt.

5.11.6 Including and Excluding Deltas

You can explicitly define which deltas you wish to include and which you wish to exclude when creating a g-file, by using the –i and –x options of the **get** command.

The –i option causes the command to apply the given deltas when constructing a version. The –x option causes the command to ignore the given deltas when constructing a version. Both options must be followed by one or more SIDs. If multiple SIDs are given they must be separated by commas (,). A range of SIDs may be given by separating two SIDs with a hyphen (–). For example, the command

```
get –i1.2,1.3 s.demo.c
```

causes deltas 1.2 and 1.3 to be used to construct the g-file. The command

```
get –x1.2–1.4 s.demo.c
```

causes deltas 1.2 through 1.4 to be ignored when constructing the file.

The –i option is useful if you wish to automatically apply changes to a version while retrieving it for editing. For example, the command

```
get –e –i4.1 –r3.3  s.demo.c
```

retrieves version 3.3 for editing. When the file is retrieved, the changes in delta 4.1 are automatically applied to it, making the g-file the same as if version 3.3 had been edited by hand using the changes in delta 4.1. These changes can be saved immediately by issuing a **delta** command. No editing is required.

The –x option is useful if you wish to remove changes performed on a given version. For example, the command

```
get –e –x1.5 –r1.6 s.demo.c
```

retrieves version 1.6 for editing. When the file is retrieved, the changes in delta 1.5 are automatically left out of it, making the g-file the same as if version 1.4 had been changed according to delta 1.6 (with no intervening delta 1.5). These changes can be saved immediately by issuing a **delta** command. No editing is required.

When deltas are included or excluded using the –i and –x options, **get** compares them with the deltas that are normally used in constructing the given version. If two deltas attempt to change the same line of the retrieved file, the command displays a warning message. The message shows the range of lines in which the problem may exist. Corrective action, if required, is the responsibility of the user.

5.11.7 Listing the Deltas of a Version

You can create a table showing the deltas required to create a given version by using the –l option. This option causes the **get** command to create an l-file which contains the SIDs of all deltas used to create the given version.

The option is typically used to create a history of a given version's development. For example, the command

```
get –l s.demo.c
```

creates a file named *l.demo.c* containing the deltas required to create the most recent version of *demo.c*.

You can display the list of deltas required to create a version by using the –lp option. The option performs the same function as the –l options except it copies the list to the standard output file. For example, the command

```
get  –lp  –r2.3  s.demo.c
```

copies the list of deltas required to create version 2.3 of *demo.c* to the standard

output.

Note that the –l option may be combined with the –**g** option to create a list of deltas without retrieving the actual version.

5.11.8 Mapping Lines to Deltas

You can map each line in a given version to its corresponding delta by using the –**m** option of the **get** command. This option causes each line in a g-file to be preceded by the SID of the delta that caused that line to be inserted. The SID is separated from the beginning of the line by a tab character. The –**m** option is typically used to review the history of each line in a given version.

5.11.9 Naming Lines

You can name each line in a given version with the current module name (i.e., the value of the %**M**% keyword) by using the –**n** option of the **get** command. This option causes each line of the retrieved file to be preceded by the value of the %**M**% keyword and a tab character.

The –**n** option is typically used to indicate that a given line is from the given file. When both the –**m** and –**n** options are specified, each line begins with the %**M**% keyword.

5.11.10 Displaying a List of Differences

You can display a detailed list of the differences between a new version of a file and the previous version by using the –**p** option of the **delta** command. This option causes the command to display the differences, in a format similar to the output of the XENIX **diff** command.

5.11.11 Displaying File Information

You can display information about a given version by using the –**g** option of the **get** command. This option suppresses the actual retrieval of a version and causes only the information about the version, such as the SID and size, to be displayed.

The –**g** option is often used with the –**r** option to check for the existence of a given version. For example, the command

```
get –g –r4.3 s.demo.c
```

displays information about version 4.3 in the s-file *s.demo.c*. If the version does not exist, the command displays an error message.

5.11.12 Removing a Delta

You can remove a delta from an s-file by using the **rmdel** command. The command has the form

rmdel –r*SID* *s.filename*

where –r*SID* gives the SID of the delta to be removed, and *s.filename* is the name of the s-file from which the delta is to be removed. The delta must be the most recently created delta in the s-file. Furthermore, the user must have write permission in the directory containing the s-file, and must either own the s-file or be the user who created the delta.

For example, the command

rmdel –r2.3 s.demo.c

removes delta 2.3 from the s-file *s.demo.c*.

The **rmdel** command will refuse to remove a protected delta, that is, a delta whose release number is below the current floor value, above the current ceiling value, or equal to a current locked value (see the section "Protecting S-files" given earlier in this chapter). The command will also refuse to remove a delta which is currently being edited.

The **rmdel** command should be reserved for those cases in which incorrect, global changes were made to an s-file.

Note that **rmdel** changes the type indicator of the given delta from "D" to "R". A type indicator defines the type of delta. Type indicators are described in full in the section *delta*(CP) in the XENIX *Reference Manual*.

5.11.13 Searching for Strings

You can search for strings in files created from an s-file by using the **what** command. This command searches for the symbol #(@) (the current value of the %Z% key word) in the given file. It then prints, on the standard output, all text immediately following the symbol, up to the next double quote ("), greater than (>), backslash (\), newline, or (non-printing) NULL character. For example, if the s-file *s.demo.c* contains the following line

char id[] = "%Z%%M%:%I%";

and the command

get –r3.4 s.prog.c

is executed, then the command

```
what prog.c
```

displays

```
prog.c:
        prog.c:3.4
```

You may also use **what** to search files that have not been created by SCCS commands.

5.11.14 Comparing SCCS Files

You can compare two versions from a given s-file by using the **sccsdiff** command. This command prints on the standard output the differences between two versions of the s-file. The command has the form

```
sccsdiff -rSID1 -rSID2 s.filename
```

where -r*SID1* and -r*SID2* give the SIDs of the versions to be compared, and *s.filename* is the name of the s-file containing the versions. The version SIDs must be given in the order in which they were created. For example, the command

```
sccsdiff -r3.4 -r5.6 s.demo.c
```

displays the differences between versions 3.4 and 5.6. The differences are displayed in a form similar to the XENIX **diff** command.

Chapter 6
Adb: A Program Debugger

6.1 Introduction 6-1

6.2 Starting and Stopping Adb 6-1
- 6.2.1 Starting With a Program File 6-1
- 6.2.2 Starting With a Core Image File 6-2
- 6.2.3 Starting Adb With Data Files 6-3
- 6.2.4 Starting With the Write Option 6-3
- 6.2.5 Starting With the Prompt Option 6-3
- 6.2.6 Leaving Adb 6-4

6.3 Displaying Instructions and Data 6-4
- 6.3.1 Forming Addresses 6-4
- 6.3.2 Forming Expressions 6-5
- 6.3.3 Choosing Data Formats 6-9
- 6.3.4 Using the = Command 6-10
- 6.3.5 Using the ? and / Commands 6-11
- 6.3.6 An Example: Simple Formatting 6-12

6.4 Debugging Program Execution 6-13
- 6.4.1 Executing a Program 6-14
- 6.4.2 Setting Breakpoints 6-15
- 6.4.3 Displaying Breakpoints 6-15
- 6.4.4 Continuing Execution 6-16
- 6.4.5 Stopping a Program with Interrupt and Quit 6-16
- 6.4.6 Single-Stepping a Program 6-16
- 6.4.7 Killing a Program 6-17
- 6.4.8 Deleting Breakpoints 6-17
- 6.4.9 Displaying the C Stack Backtrace 6-17
- 6.4.10 Displaying CPU Registers 6-18
- 6.4.11 Displaying External Variables 6-18
- 6.4.12 An Example: Tracing Multiple Functions 6-19

6.5 Using the Adb Memory Maps 6-23
- 6.5.1 Displaying the Memory Maps 6-23

6.5.2 Changing the Memory Map 6-25
6.5.3 Creating New Map Entries 6-25
6.5.4 Validating Addresses 6-26

6.6 Miscellaneous Features 6-26
6.6.1 Combining Commands on a Single Line 6-27
6.6.2 Creating Adb Scripts 6-27
6.6.3 Setting Output Width 6-28
6.6.4 Setting the Maximum Offset 6-28
6.6.5 Setting Default Input Format 6-29
6.6.6 Using XENIX Commands 6-29
6.6.7 Computing Numbers and Displaying Text 6-29
6.6.8 An Example: Directory and Inode Dumps 6-30

6.7 Patching Binary Files 6-32
6.7.1 Locating Values in a File 6-32
6.7.2 Writing to a File 6-32
6.7.3 Making Changes to Memory 6-33

6.1 Introduction

Adb is a debugging tool for C and assembly language programs. It carefully controls the execution of a program while letting you examine and modify the program's data and text areas.

This chapter explains how to use **adb**. In particular, it explains how to

- Start the debugger
- Display program instructions and data
- Run, breakpoint, and single-step a program
- Patch program files and memory

It also illustrates techniques for debugging C programs, and explains how to display information in non-ASCII data files.

6.2 Starting and Stopping Adb

Adb provides a powerful set of commands to let you examine, debug, and repair executable binary files as well as examine non-ASCII data files. To use these commands you must invoke **adb** from a shell command line and specify the file or files you wish to debug. The following sections explain how to start **adb** and describe the types of files available for debugging.

6.2.1 Starting With a Program File

You can debug any executable C or assembly language program file by typing a command line of the form

adb [*filename*]

where *filename* is the name of the program file to be debugged. **Adb** opens the file and prepares its text (instructions) and data for subsequent debugging. For example, the command

adb sample

prepares the program named "sample" for examination and execution.

Once started, **adb** normally prompts with an asterisk (*) and waits for you to type commands. If you have given the name of a file that does not exist or is in the wrong format, **adb** will display an error message first, then wait for commands. For example, if you invoke **adb** with the command

```
adb sample
```

and the file "sample" does not exist, **adb** displays the message "adb: cannot open 'sample'"

You may also start **adb** without a filename. In this case, **adb** searches for the default file *a.out* in your current working directory and prepares it for debugging. Thus, the command

```
adb
```

is the same as typing

```
adb a.out
```

Adb displays an error message and waits for a command if the *a.out* file does not exist.

6.2.2 Starting With a Core Image File

Adb also lets you examine the core image files of programs that caused fatal system errors. Core image files contain the contents of the CPU registers, stack, and memory areas of the program at the time of the error and provide a way to determine the cause of an error.

To examine a core image file with its corresponding program, you must give the name of both the core and and the program file. The command line has the form

```
adb programfile corefile
```

where *programfile* is the filename of the program that caused the error, and *corefile* is the filename of the core image file generated by the system. **Adb** then uses information from both files to provide responses to your commands.

If you do not give a core image file, **adb** searches for the default core file, named *core*, in your current working directory. If such a file is found, **adb** uses it regardless of whether or not the file belongs to the given program. You can prevent **adb** from opening this file by using the hyphen (–) in place of the core filename. For example, the command

```
adb sample –
```

prevents **adb** from searching your current working directory for a core file.

6.2.3 Starting Adb With Data Files

You can use **adb** to examine data files by giving the name of the data file in place of the program or core file. For example, to examine a data file named *outdata*, type

```
adb outdata
```

Adb opens this file and lets you examine its contents.

This method of examining files is very useful if the file contains non-ASCII data. **Adb** provides a way to look at the contents of the file in a variety of formats and structures. Note that **adb** may display a warning when you give the name of non-ASCII data file in place of a program file. This usually happens when the content of the data file is similar to a program file. Like core files, data files cannot be executed.

6.2.4 Starting With the Write Option

You can make changes and corrections in a program or data file using **adb** if you open it for writing using the **–w** option. For example, the command

```
adb –w sample
```

opens the program file *sample* for writing. You may then use **adb** commands to examine and modify this file.

Note that the **–w** option causes **adb** to create a given file if it does not already exist. The option also lets you write directly to memory after executing the given program. See the section "Patching Binary Files" later in this chapter.

6.2.5 Starting With the Prompt Option

You can define the prompt used by **adb** by using the **–p** option. The option has the form

```
–p prompt
```

where *prompt* is any combination of characters. If you use spaces, enclose the *prompt* in quotes. For example, the command

```
adb –p "Mar 10–>" sample
```

sets the prompt to "Mar 10–>". The new prompt takes the place of the default prompt (*) when **adb** begins to prompt for commands.

Make sure there is at least one space between the –p and the new prompt, otherwise **adb** will display an error message. Note that **adb** automatically supplies a space at the end of the new prompt, so you do not have to supply one.

6.2.6 Leaving Adb

You can stop **adb** and return to the system shell by using the **$q** or **$Q** commands. You can also stop the debugger by typing CNTRL-D.

You cannot stop **adb** command by pressing the INTERRUPT or QUIT keys. These keys are caught by **adb** and cause it to to wait for a new command.

6.3 Displaying Instructions and Data

Adb provides several commands for displaying the instructions and data of a given program and the data of a given data file. The commands have the form

address [, *count*] = *format*

address [, *count*] ? *format*

address [, *count*] / *format*

where *address* is a value or expression giving the location of the instruction or data item, *count* is an expression giving the number of items to be displayed, and *format* is an expression defining how to display the items. The equal sign (=), question mark (?), and slash (/) tell **adb** from what source to take the item to be displayed.

The following sections explain how to form addresses, how to choose formats, and the meaning of each of the display commands.

6.3.1 Forming Addresses

In **adb**, every address has the form

[*segment* :] *offset*

where *segment* is an expression giving the address of a specific segment of 8086/286 memory, and *offset* is an expression giving an offset from the beginning of the specified segment to the desired item. Segments and offsets are formed by combining numbers, symbols, variables, and operators. The following are some valid addresses

```
0:1
0x0bce:772
```

The *segment*: is optional. If not given, the most recently typed segment is used.

6.3.2 Forming Expressions

Expressions may contain decimal, octal, and hexadecimal integers, symbols, adb variables, register names, and a variety of arithmetic and logical operators.

Decimal, Octal, and Hexadecimal Integers

Decimal integers must begin with a nonzero decimal digit. Octal numbers must begin with a zero and may have octal digits only. Hexadecimal numbers must begin with the prefix "0x" and may contain decimal digits and the letters "a" through "f" (in both upper and lowercase). The following are valid numbers

Decimal	Octal	Hexadecimal
34	042	0x22
4090	07772	0xffa

Although decimal numbers are displayed with trailing decimal point (.), you must not use the decimal point when typing the number.

Symbols

Symbols are the names of globol variables and functions defined within the program being debugged and are equal to the address of the given variable or function. Symbols are stored in the program's symbol table and are available if the symbol table has not been stripped from the program file (see ***strip***(CP)).

In expressions, you may spell the symbol exactly as it is in the source program or as it has been stored in the symbol table. Symbols in the symbol table are no more than eight characters long and those defined in C programs are given a leading underscore (_). The following are examples of symbols.

main _main hex2bin __out_of

Note that if the spelling of any two symbols is the same (except for a leading underscore), *adb* will ignore one of the symbols and allow references only to the other. For example, if both "main" and "_main" exist in a program, then adb accesses only the first to appear in the source and ignores the other.

When you use the ? command, adb uses the symbols found in the symbol table of the program file to create symbolic addresses. Thus, the command sometimes gives a function name when displays data. This does not happen if the ? command is used for text (instructions) and the / command for data. Local variables cannot be addressed.

Adb Variables

Adb automatically creates a set of its own variables whenever you start the debugger. These variables are set to the addresses and sizes of various parts of the program file as defined below.

d	size of data
e	entry address of the program
m	execution type
n	number of segments
s	size of stack
t	size of text

Adb reads the program file to find the values for these variables. If the file does not seem to be a program file, then **adb** leaves the values undefined.

You can use the current value of an variable in an expression by preceding the variable name with an less than (<) sign. For example, the current value of the base variable "b" is

```
<b
```

You can create your own variables or change the value of an existing variable by assigning a value to a variable name with the greater than (>) sign. The assignment has the form

expression > *variable-name*

where *expression* is the value to be assigned to the variable, and *variable-name* must be a single letter. For example, the assignment

```
0x2000>b
```

assigns the hexadecimal value "0x2000" to the variable "b".

You can display the value of all currently defined **adb** variables by using the **$v** command. The command lists the variable names followed by their values in the current format. The command displays any variable whose value is not zero. If a variable also has a nonzero segment value, the variable's value is displayed as an address; otherwise it is displayed as a number.

Current Address

Adb has two special variables that keep track of the last address to be used in a command and the last address to be typed with a command. The . (dot) variable, also called the current address, contains the last address to be used in a command. The " (double quotation mark) variable contains the last address to be typed with a command. The . and " variables are usually the same except

when implied commands, such as the newline and caret (^) characters, are used. (These automatically increment and decrement ., but leave '' unchanged.)

Both the . and the '' may be used in any expression. The less than (<) sign is not required. For example, the command

```
.=
```

displays the value of the current address and

```
''=
```

displays the last address to be typed.

Register Names

Adb lets you use the current value of the CPU registers in expressions. You can give the value of the register by preceding its name with the less than (<) sign. **Adb** recognizes the following register names:

ax	register a
bx	register b
cx	register c
dx	register d
di	data index
si	stack index
bp	base pointer
fl	status flag
ip	instruction pointer
cs	code segment
ds	data segment
ss	stack segment
es	extra segment
sp	stack pointer

For example, the value of the "ax" register can be given as

```
<ax
```

Note that register names may not be used unless **adb** has been started with a *core* file or the program is currently being run under **adb** control.

Operators

You may combine integers, symbols, variables, and register names with the following operators:

Unary

~	Not
–	Negative
*	Contents of location

Binary

+	Addition
–	Subtraction
*	Multiplication
%	Integer division
&	Bitwise AND
|	Bitwise inclusive OR
^	Modulo
#	Round up to the next multiple

Unary operators have higher precedence than binary operators. All binary operators have the same precedence. Thus, the expression

```
2*3+4
```

is equal to 10 and

```
4+2*3
```

is 18.

You can change the precedence of the operations in an expression by using parentheses. For example, the expression

```
4+(2*3)
```

is equal to 10.

Note that **adb** uses 32 bit arithmetic. This means that values that exceed 2,147,483,647 (decimal) are displayed as negative values.

Note that the unary * operator treats the given address as a pointer. An expression using this operator resolves to the value pointed to by that pointer. For example, the expression

```
*0x1234
```

is equal to the value at the address "0x1234", whereas

```
0x1234
```

is just equal to "0x1234".

6.3.3 Choosing Data Formats

A format is a letter or character that defines how data is to be displayed. The following are the most commonly used formats:

Letter	Format
o	1 word in octal
d	1 word in decimal
D	2 words in decimal
x	1 word in hexadecimal
X	2 words in hexadecimal
u	1 word as an unsigned integer
f	2 words in floating point
F	4 words in floating point
c	1 byte as a character
s	a null terminated character string
i	machine instruction
b	1 byte in octal
a	the current symbolic address
A	the current absolute address
n	a newline
r	a blank space
t	a horizontal tab

A format may be used by itself or combined with other formats to present a combination of data in different forms.

The d,o,x, and u formats may be used to display **int** type variables; D and X to display **long** variables or 32-bit values. The f and F formats may be used to display single and double precision floating point numbers. The c format displays **char** type variables and s is for arrays of **char** that end with a null character (null terminated strings).

The i format displays machine instructions in 8086/286 mnemonics. The b format displays individual bytes and is useful for display data associated with instructions or the high or low bytes of registers.

The a,r, and n formats are usually combined with other formats to make the display more readable. For example, the format

```
ia
```

causes the current address to be displayed after each instruction.

You may precede each format with a count of the number of times you wish it to be repeated. For example the format

```
4c
```

displays four ASCII characters.

It is possible to combine format requests to provide elaborate displays. For example, the command

```
<b,-1/4o4^8Cn
```

displays four octal words followed by their ASCII interpretation from the data space of the core image file. In this example, the display starts at the address "<b", the base address of the program's data. The display continues until the end-of-the-file since the negative count "-1" cause an indefinite execution of the command until an error condition such as the end of the file occurs. In the format, "4o" displays the next four words (16-bit values) as octal numbers. The "4^" then moves the current address back to the beginning of these four words and "*C" redisplays them as eight ASCII characters. Finally, "n" sends a newline character to the terminal. The C format causes values to be displayed as ASCII characters if they are in the range 32 to 126. If the value is in the range 0 to 31, it is displayed as an "at" sign (@) followed by a lowercase letter. For example, the value 0 is displayed as "@a". The "at" sign itself is displayed as a double at sign "@@".

6.3.4 Using the = Command

The = command displays a given address in a given format. The command is used primarily to display instruction and data addresses in simpler form, or to display the results of arithmetic expressions. For example, the command

```
main=A
```

displays the absolute address of the symbol "main" (giving the segment and offset) and the command

```
<b+0x2000=D
```

displays (in decimal) the sum of the variable "b" and the hexadecimal value "0x2000".

If a count is given, the same value is repeated that number of times. For example, the command

```
main,2=x
```

displays the value of "main" twice.

If no address is given, the current address is used instead. This is the same as the command

```
.=
```

If no format is given, the previous format given for this command is used. For example, in the following sequence of commands both "main" and "start" are displayed in hexadecimal.

```
main=x
start=
```

6.3.5 Using the ? and / Commands

You can display the contents of a text or data segment with the ? and / commands. The commands have the form

[*address*] [, *count*] ? [*format*]

[*address*] [, *count*] / [*format*]

where *address* is an address with the given segment, *count* is the number of items you wish to display, and *format* is the format of the items you wish to display.

The ? command is typically used to display instructions in a given text segment. For example, the command

```
main,5?ia
```

displays five instructions starting at the address "main" and the address of each instruction is displayed immediately before it. The command

```
main,5?i
```

displays the instructions but no addresses other than the starting address.

The / command is typically used to check the values of variables in a program, especially variables for which no name exists in the program's symbol table. For example, the command

```
<bp-4?x
```

displays the value (in hexadecimal) of a local variable. Local variables are generally at some offset from the address pointed to by the **bp** register.

6.3.6 An Example: Simple Formatting

This example illustrates how to combine formats in ? or / commands to display different types of values whene stored together in the same program. The program to be examined has the following source statements.

```
char    str1[]   = "This is a character string" ;
int     one      = 1 ;
int     number   = 456 ;
long    lnum     = 1234 ;
float   fpt      = 1.25 ;
char    str2[]   = "This is the second character string" ;

main()
{
        one = 2;
}
```

The program is compiled and stored in a file named *sample*.

To start the session, type

```
adb sample
```

You can display the value of each individual variable by giving its name and corresponding format in a / command. For example, the command

```
str1/s
```

displays the contents of "str1" as a string

```
_str1:    This is a character string
```

and the command

```
number/d
```

displays the contents of "number" as a decimal integer

```
_number:          456.
```

You may choose to view a variable in a variety of formats. For example, you can display the **long** variable "lnum" as a 4-byte decimal, octal, and hexadecimal number by using the commands

```
lnum/D
_lnum:  1234
lnum/O
_lnum:  02322
lnum/X
_lnum:  0x4D2
```

You can also examine all variables as a whole. For example, if you wish to see them all in hexadecimal, type

```
str1,5/8x
```

This command displays eight hexadecimal values on a line and continues for five lines.

Since the data contains a combination of numeric and string values, it is worthwhile to display each value as both a number and a character to see where the actual strings are located. You can do this with one command by typing

```
str1,5/4x4^8Cn
```

In this case, the command displays four values in hexadecimal, then the same values as eight ASCII characters. The caret (^) is used four times just before displaying the characters to set the current address back to the starting address for that line.

To make the display easier to read, you can insert a tab between the values and characters and give an address for each line by typing

```
str1,5/4x4^8t8Cna
```

6.4 Debugging Program Execution

Adb provides a variety of commands to control the execution of programs being debugged. The following sections explain how to use these commands as well as how to display the contents of memory and registers.

Note that C does not generate statement labels for programs. This means it is not possible to refer to individual C statements when using the debugger. In order to use execution commands effectively, you must be familiar with the instructions generated by the C compiler and how they relate to individual C statements. One useful technique is to create an assembly language listing of your C program before using adb, then refer to the listing as you use the debugger. To create an assembly language listing, use the –S option of the cc command (see Chapter 2, "Cc: a C Compiler").

6.4.1 Executing a Program

You can execute a program by using the :r or :R commands. The commands have the form

[*address*] [,*count*] :r [*arguments*]

[*address*] [,*count*] :R [*arguments*]

where *address* gives the address at which to start execution, *count* is the number of breakpoints you wish to skip before one is taken, and *arguments* are the command line arguments, such as filenames and options, you wish to pass to the program.

If no *address* is given, then the start of the program is used. Thus, to execute the program from the beginning type

:r

If a *count* is given, **adb** will ignore all breakpoints until the given number have been encountered. For example, the command

,5:r

causes **adb** to skip the first 5 breakpoints.

If arguments are given, they must be separated by at least one space each. The arguments are passed to the program in the same way the system shell passes command line arguments to a program. You may use the shell redirection symbols if you wish.

The :R command passes the command arguments through the shell before starting program execution. This means you can use shell metacharacters in the arguments to refer to multiple files or other input values. The shell expands arguments containing metacharacters before passing them on to the program.

The command is especially useful if the program expects multiple filenames. For example, the command

:R [a-z]*.s

passes the argument "[a-z]*.s" to the shell where it is expanded to a list of the corresponding filenames before being passed to the program.

The :r and :R commands remove the contents of all registers and destroy the current stack before starting the program. This kills any previous copy of the program you may have been running.

6.4.2 Setting Breakpoints

You can set a breakpoint in a program by using the **:br** command. Breakpoints cause execution of the program to stop when it reaches the specified address. Control then returns to **adb**. The command has the form

address [, *count*] :br *command*

where *address* must be a valid instruction address, *count* is a count of the number of times you wish the breakpoint to be skipped before it causes the program to stop, and *command* is the **adb** command you wish to execute when the breakpoint is taken.

Breakpoints are typically set to stop program execution at a specific place in the program, such as the beginning of a function, so that the contents of registers and memory can be examined. For example, the command

main:br

sets a breakpoint at the start of the function named "main". The breakpoint is taken just as control enters the function and before the function's stack frame is created.

A breakpoint with a count is typically used within a function which is called several times during execution of a program, or within the instructions that correspond to a **for** or **while** statement. Such a breakpoint allows the program to continue to execute until the given function or instructions have been executed the specified number of times. For example, the command

light,5:br

sets a breakpoint at the fifth invocation of the function "light". The breakpoint does not stop the function until it has been called at least five times.

Note that no more than 16 breakpoints at a time are allowed.

6.4.3 Displaying Breakpoints

You can display the location and count of each currently defined breakpoint by using the **$b** command. The command displays a list of the breakpoints given by address. If the breakpoint has a count and/or a command, these are given as well.

The **$b** command is useful if you have creating several breakpoints in your program.

6.4.4 Continuing Execution

You can continue the execution of a program after it has been stopped by a breakpoint by using the :co command. The command has the form

[*address*] [,*count*] :co [*signal*]

where *address* is the address of the instruction at which you wish to continue execution, *count* is the number of breakpoints you wish to ignore, and *signal* is the number of the signal to send to the program (see *signal*(S) in the XENIX *Reference Manual*).

If no *address* is given, the program starts at the next instruction after the breakpoint. If a *count* is given, adb ignores the first *count* breakpoints.

6.4.5 Stopping a Program with Interrupt and Quit

You can stop execution of a program at any time by pressing the INTERRUPT or QUIT keys. These keys stop the current program and return control to **adb**, The key are especially useful for programs that have infinite loops or other program errors.

Note that whenever you press the INTERRUPT or QUIT key to stop a program, **adb** automatically saves the signal and passes it to the program if you start it again by using the :co command. This is very useful if you wish to test a program that uses these signals as part of its processing.

If you wish to continue execution of the program but do not wish to send the signals, type

:co 0

The command argument "0" prevents a pending signal from being sent to the program.

6.4.6 Single-Stepping a Program

You can single-step a program, i.e., execute it one instruction at a time, by using the :s command. The command executes an instruction and returns control to **adb**. The command has the form

[*address*] [, *count*] :s

where *address* must be the address of the instruction you wish to execute, and *count* is the number of times you wish to repeat the command.

If no *address* is given, **adb** uses the current address. If a *count* is given, **adb** continues to execute each successive instruction until *count* instructions have been executed. For example, the command

```
main,5:s
```

executes the first 5 instructions in the function *main*.

6.4.7 Killing a Program

You can kill the program you are debugging by using the **:k** command. The command kills the process created for the program and returns control to **adb**. The command is typically used to clear the current contents of the CPU registers and stack and begin the program again.

6.4.8 Deleting Breakpoints

You can delete a breakpoint from a program by using the **:dl** command. The command has the form

```
address :dl
```

where *address* is the address of the breakpoint you wish to delete.

The **:dl** command is typically used to delete breakpoints you no longer wish to use. The following command deletes the breakpoint set at the start of the function "main".

```
main:dl
```

6.4.9 Displaying the C Stack Backtrace

You can trace the path of all active functions by using the **$c** command. The command lists the names of all functions which have been called and have not yet returned control, as well as the address from which each function was called and the arguments passed to it.

For example, the command

```
$c
```

displays a backtrace of the C language functions called.

By default, the **$c** command displays all calls. If you wish to display just a few, you must supply a count of the number of calls you wish to see. For example, the command

```
,25$c
```

displays upto 25 calls in the current call path.

Note that function calls and arguments are put on the stack after the function has been called. If you put breakpoints at the entry point to a function, the function will not appear in the list generated by the **$c** command. You can remedy this problem by placing breakpoints a few instructions into the function.

6.4.10 Displaying CPU Registers

You can display the contents of all CPU registers by using the **$r** command. The command displays the name and contents of each register in the CPU as well as the current value of the program counter and the instruction at the current address. The display has the form

```
ax      0x0             fl      0x0
bx      0x0             ip      0x0
cx      0x0             cs      0x0
dx      0x0             ds      0x0
di      0x0             ss      0x0
si      0x0             es      0x0
sp      0x0             sp      0x0
0:0:    addb    al,bl
```

The value of each register is given in the current default format.

6.4.11 Displaying External Variables

You can display the values of all external variables in the program by using the $e command. External variables are the variables in your program that have global scope or have been defined outside of any function. This may include variables that have been defined in library routines used by your program.

The **$e** command is useful whenever you need a list of the names for all available variables or to quickly summarize their values. The command displays one name on each line with the variable's value (if any) on the same line.

The display has the form

```
fac:        0
_errno:     0
_end:       0
__sobuf: 0
_obuf:      0
__lastbu:   0406
__sibuf:    0
__stkmax:   0
Iscadr:     02
__iob:      01664
_edata:     0
```

6.4.12 An Example: Tracing Multiple Functions

The following example illustrates how to execute a program under **adb** control. In particular, it shows how to set breakpoints, start the program, and examine registers and memory. The program to be examined has the following source statements.

```
int     fcnt,gcnt,hcnt;
h(x,y)
{
        int hi; register int hr;
        hi = x+1;
        hr = x-y+1;
        hcnt++ ;
        hj:
        f(hr,hi);
}

g(p,q)
{
        int gi; register int gr;
        gi = q-p;
        gr = q-p+1;
        gcnt++ ;
        gj:
        h(gr,gi);
}

f(a,b)
{
        int fi; register int fr;
        fi = a+2*b;
        fr = a+b;
        fcnt++ ;
        fj:
        g(fr,fi);
}

main()
{
        f(1,1);
}
```

The program is compiled and stored in the file named *sample*. To start the session, type

```
adb sample
```

This starts **adb** and opens the corresponding program file. There is no core image file.

The first step is to set breakpoints at the beginning of each function. You can do this with the **:br** command. For example, to set a breakpoint at the start of the function "f", type

```
f:br
```

You can use similar commands for the "g" and "h" functions. Once you can created the breakpoints you can display their locations by typing

```
$b
```

This command lists the address, optional count, and optional command associated with each breakpoint. In this case, the command displays

```
breakpoints
count   bkpt            command
1       _f
1       _g
1       _h
```

The next step is to display the first five instructions in the "f" function. Type

```
f,5?ia
```

This command displays five instructions, each preceded by its symbolic address. The instructions in 8086/286 mnemonics are

```
_f:     push    bp
_f+1.:  mov     bp,sp
_f+3.:  push    di
_f+4.:  push    si
_f+5.:  call    chkstk
_f+8.:
```

You can display five instructions in "g" without their addresses by typing

```
g,5?i
```

In this case, the display is

```
_g:     push    bp
        mov     bp,sp
        push    di
        push    si
        call    chkstk
```

To start program execution, type

```
:r
```

Adb displays the message

```
sample: running
```

and begins to execute. As soon as adb encounters the first breakpoint (at the beginning of the "f" function), it stops execution and displays the message

```
breakpoint      _f:     push    bp
```

Since execution to this point caused no errors, you can remove the first breakpoint by typing

```
f:dl
```

and continue the program by typing

```
:co
```

Adb displays the message

```
sample: running
```

and starts the program at the next instruction. Execution continues until the next breakpoint where adb displays the message

```
breakpoint      _g:     push    bp
```

You can now trace the path of execution by typing

```
$c
```

The commands shows that only two functions are active: "main" and "f".

```
_f (1.,1.) from _main+6.
_main   (1.,470.)  from _start+114.
```

Although the breakpoint has been set at the start of function "g" it will not be listed in the backtrace until its first few instructions have been executed. To execute these instructions, type

```
,5:s
```

Adb single-steps the first five instructions. Now you can list the backtrace again. Type

```
$c
```

This time the list shows three active functions:

```
_g (2.,3.) from _f+48.
_f (1.,1.) from _main+6.
_main (1.,470.)     from _start+114.
```

You can display the contents of the integer variable "fcnt" by typing

```
fcnt/d
```

This command displays the value of "fcnt" found in memory. The number should be "1".

You can continue execution of the program and skip the first 10 breakpoints by typing

```
,10:co
```

Adb starts the program and display the running message again. It does not stop the program until exactly ten breakpoints have been encountered. It displays the message

```
breakpoint      _g:     push    bp
```

To show that these breakpoints have been skipped, you can display the backtrace again using **$c**.

```
_f (2., 11.)        from _h+46:
_h (10., 9.)        from _g+48:
_g (11., 20.)       from _f+48:
_f (2., 9.)         from _h+46:
_h (8., 7.)         from _g+48:
_g (9., 16.)        from _f+48:
_f (2., 7.)         from _h+46:
_h (6., 5.)         from _g+48:
_g (7., 12.)        from _f+48:
_f (2., 5.)         from _h+46:
_h (4., 3.)         from _g+48:
_g (5., 8.)         from _f+48:
_f (2., 3.)         from _h+46:
_h (2., 1.)         from _g+48:
```

6.5 Using the Adb Memory Maps

Adb prepares a set of maps for the text and data segments in your program and uses these maps to access items that you request for display. The following sections describe how to view these maps and how they are used to access the text and data segments.

6.5.1 Displaying the Memory Maps

You can display the contents of the memory maps by using the **$m** command. The command has the form

```
$m [ segment ]
```

where *segement* is the number of a segment used in the program.

The command displays the maps for all segments in the program using information taken from either the program and core files or directly from memory.

If you have started **adb** but have not executed the program, the **$m** command display has the form

```
Text Segments
Seg #    File Pos        Phys Size       'sample' - File
63.      32.             2048.
71.      2080.           656.

Data Segments
Seg #    File Pos        Phys Size       'core' - File
39.      2736.           242.
```

Each entry gives the segment number, file position, and physical size of a segment. The segment number is the starting address of the segment. The file position is the offset from the start of the file to the contents of the segment. The physical size is the number of bytes the segment occupies in the program or core file. The filenames to the right of the display are the program and core filenames.

If you have executed the program, the command display has the form

```
Text Segments
Seg #    File Pos        Vir Size  'sample' - Memory
63.      32.             2048.
71.      2080.           656.

Data Segments
Seg #    File Pos        Vir Size  'sample' - Memory
39.      2736.           456.
```

where virtual size is the number of bytes the segment occupies in memory. This size is sometimes different than the size of the segment in the file and will often change as you execute the program. This is due to expansion of the stack or allocation of additional memory during program execution. The filenames to the right always name program file. The file position value is ignored.

If you give a segment number with the command, **adb** displays information only about that segment. For example, the command

```
$m 63
```

displays a map for segment 63 only. The display has the form

```
Segment #= 63.
Type= Text
File position= 32.
Physical Size= 2048.
```

6.5.2 Changing the Memory Map

You can change the values of a memory map by using the **?m** and **/m** commands. These commands assign specified values to the corresponding map entries. The commands have the form

?m *segment-number file-position size*

and

/m *segment-number file-position size*

where *segment-number* gives the number of the segment map you wish to change, *file-position* gives the offset in the file to the beginning of the given address, and *size* gives the segment size in bytes. The **?m** assigns values to a text segment entry; **/m** to a data segment entry.

For example, the following command changes the file position for segment 63 in the text map to 0x2000:

```
?m 63 0x2000
```

The command

```
/m 39 0x0
```

changes the file position for segment 39 in the data map to 0.

6.5.3 Creating New Map Entries

You can create new segment maps and add them to your memory map by using the **?M** and **/M** commands. Unlike **?m** and **/m**, these commands create a new map instead of changing an existing one. These commands have the form

?M *segment-number file-position size*

and

/M *segment-number file-position size*

where *segment-number* gives the number of the segment map you wish to create, *file-position* gives the offset in the file to the beginning of the given

address, and *size* gives the segment size in bytes. The **?M** command creates a text segment entry; **/M** creates a data segment entry. The segment number must be unique. You cannot create a new map entry that has the same number as an existing one.

The **?M** and **/M** commands are especially useful if you wish to access segments that are otherwise allocated to your program. For example, the command

```
?M 71 0 2504
```

creates a text segment entry for segment "71" whose size is "2504" bytes.

6.5.4 Validating Addresses

Whenever you use an address in a command, **adb** checks the address to make sure it is valid. **Adb** uses the segment number, file position, and size values in each map entry to validate the addresses. If an address is correct, **adb** carries out the command; otherwise, it displays an error message.

The first step **adb** takes when validating an address is to check the segment value to make sure it belongs to the appropriate map. Segments used with the **?** command must appear in the text segments map; segments used with the **/** command must appear in the data segments map. If the value does not belong to the map, **adb** displays a bad segment error.

The next step is to check the offset to see if it is in range. The offset must be within the range

```
0 <= offset <= segment-size
```

If it is not in this range, **adb** displays a bad address error.

If **adb** is currently accessing memory, the validating segment and offset are used to access a memory location and no other processing takes place. If **adb** is accessing files, it computes an effective file address

```
effective-file-address = offset + file-position
```

then uses this effective address to read from the corresponding file.

6.6 Miscellaneous Features

The following sections explain how to use a number of useful commands and features of **adb**.

6.6.1 Combining Commands on a Single Line

You can give more than one command on a line by separating the commands with a semicolon (;). The commands are performed one at a time, starting at the left. Changes to the current address and format are carried to the next command. If an error occurs, the remaining commands are ignored.

One typical combination is to place a **?** command after a **l** command. For example, the commands

```
?l 'Th'; ?s
```

search for and display a string that begins with the characters "Th".

6.6.2 Creating Adb Scripts

You can direct **adb** to read commands from a text file instead of the keyboard by redirecting **adb**'s standard input file at invocation. To redirect the standard input, use the standard redirection symbol < and supply a filename. For example, to read commands from the file *script*, type

```
adb sample <script
```

The file you supply must contain valid **adb** commands. Such files are called script files and can be used with any invocation of the debugger.

Reading commands from a script file is very convenient when you wish to use the same set of commands on several different object files. Scripts are typically used to display the contents of core files after a program error. For example, a file containing the following commands can be used to display most of the relevant information about a program error:

```
120$w
4095$s
$v
=3n
$m
=3n"C Stack Backtrace"
$C
=3n"C External Variables"
$e
=3n"Registers"
$r
0$s
=3n"Data Segment"
<b,-1/8xna
```

6.6.3 Setting Output Width

You can set the maximum width (in characters) of each line of output created by **adb** by using the **$w** command. The command has the form

n$w

where *n* is an integer number giving the width in characters of the display. You may give any width that is convenient for your given terminal or display device. The default width when **adb** is first invoked is 80 characters.

The command is typically used when redirecting output to a lineprinter or special terminal. For example, the command

120$w

sets the display width to 120 characters, a common maximum width for lineprinters.

6.6.4 Setting the Maximum Offset

Adb normally displays memory and file addresses as the sum of a symbol and an offset. This helps associate the instructions and data you are viewing with a given function or variable. When first invoked, **adb** sets the maximum offset to 255. This means instructions or data that are no more than 255 bytes from the start of the function or variable are given symbolic addresses. Instructions or data beyond this point are given numeric addresses.

In many programs, the size of a function or variable is actually larger than 255 bytes. For this reason **adb** lets you change the maximum offset to accomodate larger programs. You can change the maximum offset by using the **$s** command. The command has the form

n$s

where *n* is an integer giving the new offset. For example, the command

4095$s

increases the maximum possible offset to 4095. All instructions and data that are no more than 4095 bytes away are given symbolic addresses.

Note that you can disable all symbolic addressing by setting the maximum offset to zero. All addresses will be given numeric values instead.

6.6.5 Setting Default Input Format

You can set the default format for numbers used in commands with the **$d** (decimal), **$o** (octal), and **$x** (hexadecimal) commands. The default format tells **adb** how to interpret numbers that do not begin with "0" or "0x" and how to display numbers when no specific format is given.

The commands are useful if you wish to work with a combination of decimal, octal, and hexadecimal numbers. For example, if you use

```
$x
```

you may give addresses in hexadecimal without prepending each address with "0x". Furthermore, **adb** displays all numbers in hexadecimal except those specifically requested to be in some other format.

When you first start **adb**, the default format is decimal. You may change this at any time and restore it as necessary using the **$d** command.

6.6.6 Using XENIX Commands

You can execute XENIX commands without leaving **adb** by using the **adb** escape command **!**. The escape command has the form

```
! command
```

where *command* is the XENIX command you wish to execute. The command must have any required arguments. **Adb** passes this command to the system shell which executes it. When finished, the shell returns control to **adb**.

For example, to display the date type

```
! date
```

The system displays the date at your terminal and restores control **adb**.

6.6.7 Computing Numbers and Displaying Text

You can perform arithmetic calculations while in **adb** by using the **=** command. The command directs **adb** to display the value of an expression in a given format.

The command is often used to convert numbers in one base to another, to double check the arithmetic performed by a program, and to display complex addresses in easier form. For example, the command

```
0x2a=d
```

displays the hexadecimal number "0x2a" as the decimal number 42 but

```
0x2a=c
```

displays it as the ASCII character "*". Expressions in a command may have any combination of symbols and operators. For example, the command

```
<d0-12*<d1+<b+5=X
```

computes a value using the contents of the d0 and d1 registers and the **adb** variable "b". You may also compute the value of external symbols as in the command

```
main+5=X
```

This is helpful if you wish to check the hexadecimal value of an external symbol address.

Note that the = command can also be used to display literal strings at your terminal. This is especially useful in **adb** scripts where you may wish to display comments about the script as it performs its commands. For example, the command

```
=3n"C Stack Backtrace"
```

spaces three lines, then prints the message "C Stack Backtrace" on the terminal.

6.6.8 An Example: Directory and Inode Dumps

This example illustrates how to create **adb** scripts to display the contents of a directory file and the inode map of a XENIX file system. The directory file is assumed to be named *dir* and contains a variety of files. The XENIX file system is assumed to be associated with the device file */dev/src* and has the necessary permissions to be read by the user.

To display a directory file, you must create an appropriate script, then start **adb** with the name of the directory, redirecting its input to the script.

First, you can create a script file named *script*. A directory file normally contains one or more entries. Each entry consists of an unsigned "inumber" and a 14 character filename. You can display this information by adding the command

```
0,-1?ut14cn
```

to the script file. This command displays one entry for each line, separating the

number and filename with a tab. The display continues to the end of the file. If you place the command

```
="inumber"8t"Name"
```

at the beginning of the script, **adb** will display the strings as headings for each column of numbers.

Once you have the script file, type

```
adb dir - <script
```

(The hyphen (-) is used to prevent **adb** from attempting to open a core file.) **Adb** reads the commands from the script and the resulting display has the form

```
inumber name
652     .
82      ..
5971    cap.c
5323    cap
0       pp
```

To display the inode table of a file system, you must create a new script, then start **adb** with the filename of the device associated with the file system (e.g., the hard disk drive).

The inode table of a file system has a very complex structure. Each entry contains: a word value for the file's status flags; a byte value for the number links; two byte values for the user and group IDs; a byte and word value for the size; eight word values for the location on disk of the file's blocks; and two word values for the creation and modification dates. The inode table starts at the address "**02000**". You can display the first entry by typing

```
02000,-1?on3bnbrdn8un2Y2na
```

Several newlines are inserted within the display to make it easier to read.

To use the script on the inode table of */dev/src*, type

```
adb /dev/src - <script
```

(Again, the hypen (-) is used to prevent an unwanted core file.) Each entry in the display has the form

```
02000:  073145
        0163        0164        0141
        0162        10356
        28770       8236        25956   27766   25455   8236    25956   25206
        1976 Feb 5 08:34:56     1975 Dec 28 10:55:15
```

6.7 Patching Binary Files

You can make corrections or changes to any file, including executable binary files, by using the **w** and **W** commands and invoking **adb** with the **–w** option. The following sections describe how to locate and change values in a file.

6.7.1 Locating Values in a File

You can locate specific values within a file by using the **l** and **L** commands. The commands have the form

[*address*] ?l *value*

where *address* is the address at which to start the search, and *value* is the value (given as an expression) to be located. The **l** command searches for 2 byte values; **L** for 4 bytes.

The

?l

commands starts the search at the current address and continues until the first match or the end of the file. If the value is found, the current address is set to that value's address. For example, the command

?l 'Th'

searches for the first occurrence of the string value "Th". If the value is found at "main+210" the current address is set to that address.

6.7.2 Writing to a File

You can write to a file by using the **w** and **W** commands. The commands have the form

[*address*] ?w *value*

where *address* is the address of the value you wish to change, and *value* is the new value. The **w** command writes 2 byte values; **W** writes 4 bytes. For example, the following commands change the word "This" to "The ".

```
?l 'Th'
?W 'The'
```

Note that **W** is used to change all four characters.

6.7.3 Making Changes to Memory

You can also make changes to memory whenever a program has been executed. If you have used an **:r** command with a breakpoint to start program execution, subsequent **w** commands cause **adb** to write to the program in memory rather than the file. This is useful if you wish to make changes to a program's data as it runs, for example, to temporarily change the value of program flags or constants.

Chapter 7
As: An Assembler

7.1 Introduction 7-1

7.2 Command Usage 7-1

7.3 Lexical Conventions 7-2
- 7.3.1 Identifiers 7-2
- 7.3.2 Constants 7-2
- 7.3.3 Whitespace 7-2
- 7.3.4 Comments 7-2

7.4 Assembly Segments 7-3
- 7.4.1 Text, Data, and Bss Segments 7-3
- 7.4.2 The Location Counter 7-4

7.5 Statements 7-4
- 7.5.1 Labels 7-4
- 7.5.2 Null Statements 7-5
- 7.5.3 Expression Statements 7-5
- 7.5.4 Assignment Statements 7-5
- 7.5.5 Keyword Statements 7-6

7.6 Expressions 7-6
- 7.6.1 Expression Operators 7-6
- 7.6.2 Types 7-6
- 7.6.3 Type Propagation in Expressions 7-7

7.7 Assembler Directives 7-8
- 7.7.1 Even Directive 7-9
- 7.7.2 Floating Point Directive 7-9
- 7.7.3 Global Directive 7-9
- 7.7.4 Segment Directives 7-10
- 7.7.5 Common Directive 7-10
- 7.7.6 Insert Directive 7-10
- 7.7.7 ASCII Directives 7-11

7.7.8 Listing Directives 7-11
7.7.9 Block Directives 7-12
7.7.10 Initial Value Directives 7-12
7.7.11 End Directive 7-12

7.8 Machine Instructions 7-13
7.8.1 Mnemonic List 7-13
7.8.2 Byte Instructions 7-16
7.8.3 Branch Instructions 7-17
7.8.4 String Instructions 7-18
7.8.5 Intersegment Instructions 7-18
7.8.6 Input/Output Instructions 7-18
7.8.7 80286 Instructions 7-18

7.9 Addressing Modes 7-19
7.9.1 Register Operands 7-19
7.9.2 Immediate Operands 7-20
7.9.3 Direct Address Operands 7-20
7.9.4 Based Operands 7-21
7.9.5 Indexed Operands 7-21
7.9.6 Based Indexed Operands 7-22
7.9.7 Indirect Address Operands 7-22

7.10 Diagnostics 7-23

7.1 Introduction

This chapter describes the usage and input syntax of the XENIX 8086/286 assembler, as. The assembler produces relocatable object files from 8086/286 assembly language source files. Object files contain relocation information and a complete symbol table, and may be linked to other object files using the XENIX loader *ld*(CP).

As is designed to be used in those rare cases where C programs do not satisfy a programming requirement. Thus, you can combine as object files with object files produced by the XENIX C compiler, cc, to make complete programs. Note that the output format of as has been designed so that if a file contains no unresolved references to external symbols, it is executable without further processing.

This chapter does not teach assembly language programming, nor does it give a detailed description of 8086/286 operation codes. For information on these topics, you will need other references.

7.2 Command Usage

As is invoked as follows:

as [*option*] *filename* ...

where *options* is an assembler option, and *filename* is the name of the assembler source file. If the *filename* does not have the extension *.s*, B as displays a warning message before assembly ing the file. Although as has a large number of options, the most commonly used are the –l and –o options.

The –l option causes the assembler to create an an assembly listing which includes the source, the assembled (binary) code, and any assembly errors. The listing file is given the *filename*.L.

The –o option causes the output to be placed in a given file. The option has the form

–o *outfile*

where *outfile* is the name of the file to receive the assembled program. If you do not use the –o, as copies the output to the file named *filename*.o in the current directory.

For a complete description of all assembler options, see *as*(CP) is the XENIX *Reference Manual*.

7.3 Lexical Conventions

Assembler tokens include identifiers, constants, and operators.

7.3.1 Identifiers

An identifier consists of a sequence of alphanumeric characters, including period "." and underscore "_". The first character must not be numeric. By convention, the first eight characters are significant, but you can also define the number of significant characters by using the **–nl** option. Uppercase and lowercase letters are considered distinct in identifiers.

7.3.2 Constants

A hex constant consists of a slash character (/) followed by a sequence of digits and one of the letters "a", "b", "c", "d", "e", or "f" any of which may be capitalized.

A decimal constant consists simply of a sequence of digits. The constant should be representable in 15 bits, i.e., be less than 32,768.

A character constant consists of one or two characters enclosed in a single quotation mark ('). If the quotation mark is used in the constant it must be given twice.

The following are examples of constants

Decimal	Hexadecimal	Character
10	/1b	'a'
32767	/7fff	'in'

7.3.3 Whitespace

Blank and tab characters may be freely interspersed between tokens, but may not be used within tokens (except in character constants). A blank or tab is required to separate adjacent identifiers or constants not otherwise separated.

7.3.4 Comments

The vertical bar (|) introduces a comment, which extends to the end of the line on which it appears. Comments are ignored by the assembler.

7.4 Assembly Segments

As assembles instruction and data statements in three segments. Segments allow division of instructions and data into separate physical segments in memory. A location counter kepts the current address within each segment during assembly and provides reference to the current instruction and data.

7.4.1 Text, Data, and Bss Segments

Every program is divided into at most three distinct segments of assembled code and data: the text segment, the data segment, and the bss segment. Each segment is reserved for a specific type of storage and receives different treatment from the assembler and from the XENIX linker when the final program is created.

The text segment is normally reserved for instructions, but may also be used for data. Instructions in this segment are assembled, and the code is copied to the output file. Data definitions in this segment are also assembled and copied; the code is the value of the data item. The assembler does not separate the instruction and data code. If the instructions and data definitions are mixed within the source file, the resulting code is mixed within the output file.

The data segment is reserved for data and instructions that may be modified during execution. Although the instructions and data definitions in this segment are processed the same as in the text segment, the code is copied to a different part of the output file and receives different treatment from the XENIX linker.

The bss segment is reserved for uninitialized data only. Instructions or data definitions with initial values must not be given in this segment. The assembler counts the number of bytes allocated for this segment and copies this count to the output file. It does not generate code.

The text segment is implicitly defined at the start of every assembly. Thus, any instructions or data definitions given when no other segment is explicitly defined are copied to the text segment. To start a data or bss segment, you must use a **.data** or **.bss** directive. You can explicitly start the text segment with the **.text** directive (see section 7.8).

Unless otherwise specified, the first statement in the text segment is considered the program's entry point. In shared-text programs, the instructions and data in the text segment are write-protected; in nonshared-text programs, they are not. Instructions and data in the data segment are never write-protected. The bss segment is actually an extension of the data segment. It begins immediately after the data segment and is initialized to 0 at the start of program execution.

7.4.2 The Location Counter

The special symbol, "dot" (.), is the location counter. Its value at any time is the offset from the current statement to the start of the current segment. Thus, it may be used in any statement to refer to the current location.

The location counter actually has three different offsets, one for each type of segment. Only the offset of the current segment is ever accessible. The assembler increments the current offset after it processes each statement. It increments the offset by the number of bytes in the assembled code or allocated storage.

The location counter can be assigned an explicit value if desired. Its value must not be decreased. If it is explicitly increased, the assembler generates enough null bytes of code to fill that gap between the last offset and the new one.

7.5 Statements

A source program is composed of a sequence of statements. Statements are separated by newline characters. There are four kinds of statements:

— Null statements

— Expression statements

— Assignment statements

— Keyword statements

The format for most 8086/286 assembly language source statements is:

[*labelfield*] op-code [*operand-field*] [*comment*]

Any kind of statement may be preceded by one or more labels.

7.5.1 Labels

There are two kinds of labels: name labels and numeric labels. A name label consists of an identifier followed by a colon (:). The effect of a name label is to assign the current value and type of the location counter to the name. An error is indicated in pass 1 if the name is already defined; an error is indicated in pass 2 if the value assigned changes the definition of the label.

A numeric label consists of a string of digits 0 to 9 and a dollar-sign ($) followed by a colon (:). Such a label serves to define local symbols of the form

n$

where *n* is the digit of the label. The scope of the numeric label is the labeled block in which it appears. As an example, the label "9$" is defined only between the labels *label1* and *label2*:

```
label1:
9$:         .byte 0
            .
            .
            .
label2:     .word a
```

As in the case of name labels, a numeric label assigns the current value and type of dot to the symbol.

7.5.2 Null Statements

A null statement is an empty statement (which may, however, have labels and a comment). A null statement is ignored by the assembler. Common examples of null statements are empty lines or lines containing only a label.

7.5.3 Expression Statements

An expression statement consists of an arithmetic expression not beginning with a keyword. The assembler computes its value and places it in the output stream, together with the appropriate relocation bits.

7.5.4 Assignment Statements

An assignment statement consists of an identifier, an equal sign (=), and an expression. The value and type of the expression are assigned to the identifier. It is not required that the type or value be the same in pass 2 as in pass 1, nor is it an error to redefine any symbol by assignment.

Any external attribute of an expression is lost across an assignment. This means that it is not possible to declare a global symbol by assigning to it, and that it is not possible to define a symbol to be offset from a nonlocally defined global symbol.

As mentioned, it is permissible to assign to the location counter. It is required, however, that the type of the expression assigned be of the same type as dot, and it is forbidden to decrease the value of dot. In practice, the most common assignment to dot has the form

. = . + *n*

for some number *n*; this has the effect of generating *n* null bytes.

7.5.5 Keyword Statements

Keyword statements are numerically the most common type, since most machine instructions are of this sort. A keyword statement begins with one of the many predefined keywords of the assembler. The syntax of the remainder depends on the keyword. All the keywords are listed in sections 7.8 and 7.9.

7.6 Expressions

An expression is a sequence of symbols representing a value. Its constituents are identifiers, constants, and operators. Each expression has a type.

Arithmetic is two's complement. All operators have equal precedence, and expressions are evaluated strictly left to right.

7.6.1 Expression Operators

The operators are:

Operator	Description
+	Addition
-	Subtraction
*	Multiplication
/	Division
^	Modulus
&	Logical AND
		Logical NOT
>>	Right Shift
<<	Left Shift

7.6.2 Types

The assembler deals with expressions, each of which may be of a different *type*. Most types are attached to the keywords and are used to select the routine which treats that keyword. The types likely to be met explicitly are:

undefined
: Upon first encounter, each symbol is undefined. A defined symbol may become undefined if it is assigned an undefined expression.

undefined external
: A symbol which is declared **.globl** but not defined in the current

assembly is an undefined external. If such a symbol is declared, the link editor *ld*(1S) must be used to load the assembler's output with another routine that defines the undefined reference.

absolute
: An absolute symbol is defined ultimately from a constant. Its value is unaffected by any possible future applications of the link-editor to the output file.

text
: The value of a text symbol is measured with respect to the beginning of the text segment of the program. If the assembler output is link-edited, its text symbols may change in value, since the program need not be the first in the link editor's output. Most text symbols are defined by appearing as labels. At the start of an assembly, the value of dot is text **0**.

data
: The value of a data symbol is measured with respect to the origin of the data segment of a program. Like text symbols, the value of a data symbol may change during a subsequent link-editor run, since previously loaded programs may have data segments. After the first *.data* statement, the value of dot is data **0**.

bss
: The value of a bss symbol is measured from the beginning of the bss segment of a program. Like text and data symbols, the value of a bss symbol may change during a subsequent link-editor run, since previously loaded programs may have bss segments. After the first **.bss** statement, the value of dot is bss **0**.

external absolute, text, data, or bss
: Symbols declared **.globl** but defined within an assembly as absolute, text, data, or bss symbols may be used exactly as if they were not declared **.globl**. However, their value and type are available to the link editor so that the program may be loaded with others that reference these symbols.

other types
: Each keyword known to the assembler has a type that is used to select the routine which processes the associated keyword statement. The behavior of such symbols when not used as keywords is the same as if they were absolute.

7.6.3 Type Propagation in Expressions

When operands are combined by expression operators, the result has a type that depends on the types of the operands and on the operator. The rules involved are complex, but are intended to be sensible and predictable. For

purposes of expression evaluation the important types are:

undefined
absolute
text
data
bss
undefined external
other

The combination rules are as follows:

— If one of the operands is undefined, the result is undefined.

— If both operands are absolute, the result is absolute.

— If an absolute is combined with one of the other types mentioned above, the result has the other type.

— If two operands of other type are combined, the result has the numerically larger type.

— An other type combined with an explicitly discussed type other than absolute acts like an absolute.

Further rules applying to particular operators are:

+ If one operand is text-, data-, or bss-segment relocatable, or is an undefined external, the result has the postulated type and the other operand must be absolute.

– If the first operand is a relocatable text-, data-, or bss-segment symbol, the second operand may be absolute (in which case the result has the type of the first operand); or the second operand may have the same type as the first (in which case the result is absolute). If the first operand is external undefined, the second must be absolute. All other combinations are illegal.

others
It is illegal to apply these operators to any but absolute symbols.

7.7 Assembler Directives

As supports a number of assembler directives (sometimes called "Pseudo-operations"). The directives modify the location counter, define the start of program segments, generate initialized data, allocate storage space, assign the global attribute to labels or symbols, and perform a variety of other tasks.

The following sections describe the directives and illustrate their use.

7.7.1 Even Directive

.even

The **.even** directive conditionally increments the location counter. If the location counter is odd, it is incremented by one so the next statement will be assembled at a word boundary. This is useful for forcing storage allocation to be on a word boundary after a **.byte** or **.ascii** directive.

7.7.2 Floating Point Directive

.float *float*
.double *float*

The **.float** and **.double** directives accept one or more floating point numbers as operands and allocate storage for each number. A floating point number has the form

[–] *integer.fraction* [E [–] *exponent*]

where *integer* is a decimal number, *fraction* is a combination of decimal digits, and *exponent* is an decimal number.

```
.float   25.1
.double  1.03E12
.float   -34718.235E4
```

The **.float** and **.double** directives allocate a different number of bytes. The **.float** sets aside four bytes, while **.double** sets aside eight.

7.7.3 Global Directive

.globl *name* ...

The **.globl** directive makes the text or data associated with *name* globally known to all files in a program. The *name* must be explicitly defined by assignment or by appearance as a label in exactly one file. All other files that wish to access this *name* must use **.globl** to give the name global meaning; no other definition is allowed in these files. The link editor *ld*(CP) resolves all global references to *name* when the final program is created. If more than one name is given, they must be separated with commas (,).

7.7.4 Segment Directives

```
.text
.data
.bss
```

The .text, .data, and .bss directives cause the assembler to copy subsequent instruction or data code (or allocated storage) to the text, data, or bss segment, respectively. The offset of the location counter is set to the previous value for that segment and subsequent statements are processed as defined in section 7.4.

The directives may be used any number of times within a program. The offset for each segment is initially set to 0. Changing a segment causes the current offset to be saved. Restoring a segment causes the old offset to be restored. Thus, each segment is copied as a contiguous block even if the original source statements were not contiguous.

Instructions and data definitions with initial values must not be used after a .bss directive, but symbols may be defined and dot moved by assignment.

If no explicit segment directive is given in a program, code is copied to the text segment.

7.7.5 Common Directive

.comm *name* [, *expression*]

The **.comm** directive makes *name* globally known to all files of the program. If *name* also appears in an assignment or as a label in a file, then the **.comm** directive has the same effect as the **.globl** directive. In this case, any *expression* given is ignored. If *name* does not have an explicit definition, then **.comm** directs the XENIX linker to automatically allocate *expression* bytes for *name* in the bss segment of the program. These bytes appear before any bytes specifically allocated within the bss segment.

7.7.6 Insert Directive

.insrt "*filename* "

The **.insrt** directive directs the assembler to suspend processing of the current file until all statements in the given file *filename* have been read. The *filename* must be enclosed within double quotation marks. If the file cannot be opened or does not exist, the assembler displays the message

Cannot open insert file

Otherwise, it reads the contents of the file. The file may contain other **.insrt** directives; upto 10 levels of directives may be nested in this way.

The **.insrt** directive is useful for including a standard set of comments or symbol assignments at the beginning of a program, e.g., the definitions for system calls found in the file */usr/include/sys.s*. The directive is also useful for breaking up a large source program into easily manageable pieces.

7.7.7 ASCII Directives

.ascii /*string*/
.asciz /*string*/

The **.ascii** and **.asciz** directives translate the *string* into an equivalent sequence of ASCII byte values and copy these bytes to allocated storage in the current segment. The **.asciz** directive also appends a null byte to the end of the sequence.

The *string* may contain any character in the character set except a newline . If necessary, the escape sequence "\n" may be used in place of a newline. The *string* must be enclosed within slashes (/) or within any character not used in the *string*.

```
.ascii /"hello there"/
.ascii "Warning-\007\007 \n"
.asciz *abcdefg*
```

The **.asciz** directive is useful when accessing strings of undetermined length. The last null byte marks the end of the string. Also, some XENIX system calls require null terminated strings as arguments.

7.7.8 Listing Directives

.list
.nlist

The **.list** and **.nlist** directives control output to the assembler listing file created by using the **-l** option. If **.list** is given, subsequent statements are passed to the listing file as well as being processed. If **.nlist** is given, statements are processed but not passed to the listing file.

The directives may be used any number of times to turn listing on and off. This is particularly useful when certain portions of the assembly output is not necessarily desired on a printed listing.

7.7.9 Block Directives

.blkb [*expression*]
.blkw [*expression*]

The .blkb and .blkw directives reserve blocks of storage where a block contains *expression* bytes (for .blkb) or *expression* words (for .blkw). If no *expression* is given, 1 is assumed. The expression must be absolute and defined during pass 1.

Note that the statement

.=.+*expression*

may also be used to reserve blocks of storage. In this cases, the block contains *expression* bytes.

7.7.10 Initial Value Directives

.byte [*expression*] ...
.word [*expression*] ...

The .byte and .word directives reserve storage and initialize this storage to the value given by *expression*. The .byte directive reserves one byte for each *expression* and initializes that byte to the low-order of the *expression*. The .word directive reserves one word for each *expression* and initializes that word to the value of *expression*. When more than one *expression* is given, they must be separated by commas (,).

7.7.11 End Directive

.end [*expression*]

The .end directive marks the physical end of the source program. If *expression* is given, it indicates the entry point of the program, i.e., the starting point for execution. Otherwise, the entry point is taken to be the start of the text segment.

Note that inserted files that contain an .end directive terminate assembly of the entire program as well as the inserted portion.

7.8 Machine Instructions

This section presents a description of the 8086/286 instructions used by the XENIX 8086/286 assembler. This assembler is intended for both the 8086 and 80286 processors. However, it does not support instructions that are specific to the 80286 processor. See section 7.9.7 for a description of how to form these instructions using directives.

7.8.1 Mnemonic List

The following is a list of all instructions mnemonics recognized by the XENIX 8086/286 assembler. Operand order and usage for most instructions are identical to that found with other 8086/286 assemblers. Although most of the mnemonics are the same as for other 8086/286 assemblers, many are unique to the XENIX system and are not necessarily compatible with other assemblers. Instructions marked by an asterisk (*) are described in more detail in later sections.

XENIX 8086/286 Assembler Mnemonics

Mnemonic	Description	XENIX Specific
aaa	ascii adjust for addition	
aad	ascii adjust for division	
aam	ascii adjust for multiply	
aas	ascii adjust for subtraction	
adc	add with carry	
adcb	add byte with carry	*
add	add	
addb	add byte	*
and	logical AND	
andb	logical AND byte	*
beq	long branch equal	*
bge	long branch greater or equal	*
bgt	long branch greater	*
bhi	long branch on high	*
bhis	long branch high or same	*
ble	long branch less than or equal	*
blo	long branch on low	*
blos	long branch low or same	*
blt	long branch less than	*
bne	long branch not equal	*
br	long branch	*
call	intra segment call	
calli	inter segment call	*
cbw	convert byte to word	
clc	clear carry flag	
cld	clear direction flag	
cli	clear interrupt flag	

cmc	complement carry flag	
cmp	compare	
cmpb	compare byte	*
cmps	compare string	*
cmpsb	compare string	*
cwd	covert word to double word	
daa	decimal adjust for addition	
das	decimal adjust for subtraction	
dec	decrement by one	
decb	decrement byte by one	*
div	division unsigned	
divb	division unsigned byte	*
hlt	halt	
idiv	integer division	
idivb	integer division byte	*
imul	integer multiplication	
imulb	integer multiplication	*
in	input byte	*
inc	increment by one	
incb	increment byte by one	*
int	interrupt	
into	interrupt if overflow	
inw	input word	*
iret	interrupt return	
j	short jump	*
ja	short jump if above	
jae	short jump if above or equal	
jb	short jump if below	
jbe	short jump if below or equal	
jc	short jump if carry	
jcxz	short jump if CX is zero	
je	short jump on equal	
jg	short jump on greater than	
jge	short jump greater than or equal	
jl	short jump on less than	
jle	short jump on less than or equal	
jmp	jump	
jmpi	inter segment jump	*
jna	short jump not above	
jnae	short jump not above or equal	
jnb	short jump not below	
jnbe	short jump not below or equal	
jnc	short jump not carry	
jne	short jump not equal	
jng	short jump not greater	
jnge	short jump not greater or equal	
jnl	short jump not less	
jnle	short jump not less or equal	
jno	short jump not overflow	
jnp	short jump not parity	

jns	short jump not sign	
jnz	short jump not zero	
jo	short jump on overflow	
jp	short jump if parity	
jpe	short jump if parity even	
jpo	short jump if parity odd	
js	short jump if signed	
jz	short jump if zero	
lahf	load AH from flags	
lds	load pointer using DS	
lea	load effective address	
les	load pointer using ES	
lock	lock bus	
lodb	load byte string	*
lodw	load word string	*
loop	loop short label	
loope	loop if equal	
loopne	loop if not equal	
loopnz	loop is not zero	
loopz	loop if zero	
mov	move	
movb	move byte	*
movs	move word string	*
movsb	move byte string	*
mul	multiplication unsigned	
mulb	multiplication unsigned	*
neg	negate	
negb	negate	*
nop	no operation	
not	logical NOT	
notb	logical NOT	*
or	logical OR	
orb	logical OR	*
out	output byte	*
outw	output word	*
pop	pop from stack	
popf	pop flag from stack	
push	push onto stack	
pushf	push flags onto stack	
rcl	rotate left through carry	
rclb	rotate left through carry	*
rcr	rotate right through carry	
rcrb	rotate right through carry	*
rep	repeat string operation	*
repnz	repeat string operation not zero	*
repz	repeat string operation while zero	*
ret	return from procedure	
reti	return from intersegment procedure	*
rol	rotate left	
rolb	rotate left	*

ror	rotate right	
rorb	rotate right	*
sahf	store AH into flags	
sal	shift arithmetic left	
salb	shift arithmetic left	*
sar	shift arithmetic right	
sarb	shift arithmetic right	*
sbb	subtract with borrow	
sbbb	subtract with borrow	*
scab	scan byte string	*
scaw	scan word string	*
shl	shift logical left	
shlb	shift logical left	*
shr	shift logical right	
shrb	shift logical right	*
stc	set carry flag	
std	set direction flag	
sti	set interrupt enable flag	
stob	store byte string	*
stow	store word string	*
sub	subtraction	
subb	subtraction	*
test	test	
testb	test	*
wait	wait while TEST pin	
xchg	exchange	
xchgb	exchange	*
xlat	translate	*
xor	exclusive OR	
xorb	exclusive OR	*

7.8.2 Byte Instructions

The XENIX assembler extends the definition of several instruction mnemonics to include an explicit byte "b" suffix. This suffix forces the operands in the instruction to be treated as bytes when they would otherwise be treated as words. There are the following byte instructions.

adcb	imulb	rclb	shlb
addb	incb	rcrb	shrb
andb	movb	rolb	subb
cmpb	mulb	rorb	testb
decb	negb	salb	xchgb
divb	notb	sarb	xorb
idivb	orb	sbbb	

The byte instructions are especially useful when operating on memory operands defined with the .byte directive. Since the XENIX assembler does not

assign an explicit type to symbols created with the **.byte** or **.word** directives, it is impossible for the assembler to detect the size of the associated item when given in an instruction. For example, if "test_byte" and "test_word" have been defined as

```
test_byte:      .byte   1
test_word:      .word   1
```

then the statements

```
negb    test_byte
neg     test_word
```

are required to operate on these values correctly. If **neg** were applied to "test_byte", part of "test_word" would be destroyed in the operation.

7.8.3 Branch Instructions

The XENIX assembler has a new class of instructions, called branch instructions, which provide a way to test a condition and branch to an instruction address that is further than 128 bytes away. These instructions take the same kind of operand as the normal **jmp** instruction, but provide a test to see whether or not the jump should occur. The following is a list of the branch instructions.

```
beq     bhis    blt
bge     ble     bne
bgt     blo     br
bhi     blos
```

The branch instructions, when assembled, consist of two 8086/286 jump instructions. The first jump tests for the inverse of the condition specified by the branch. The second jump is the unconditional jump instruction **jmp**. If the branch condition is true, the first jump is ignored and the second jump taken. If the branch condition is false, the first jump is taken. The first jump passes control to the next statement after the second jump.

For example, the statement

```
bne     subtest
```

is equivalent to the statements

```
    je      no
    jmp     subtest
no:
```

7.8.4 String Instructions

The XENIX assembler uses a subset of the string instructions normally available for the 8086/286 processor. In particular, the assembler accepts only those string instructions which do not take operands. These are essentially the byte and word forms of the string instructions with implied destination.

To indicate this restriction, some instruction mnemonics no longer contain the "s" for string. The following is a list of the string instructions.

cmpsb	lodb	movsb	scab	stob
cmps	lodw	movs	scaw	stow

The assembler also accepts the **rep**, **repnz**, and **repz** instructions for repeating string operations. Note that these instructions must appear alone on a line; they cannot be combined on the same line with a string instructions.

7.8.5 Intersegment Instructions

The XENIX assembler has redefined call and jump instructions to create a new class of instructions, called intersegment instructions. These allow calls and unconditional jumps to locations across 8086/286 segment boundaries. (These are physical segments, not the text, data, and bss segments described earlier.) There are the following instructions.

```
calli
jmpi
reti
```

The **calli** and **jmpi** instructions can have either a locally or globally defined symbol as an operand. In this case, an appropriate segment address is provided automatically when the program is linked. If an indirect address operand is used (see section 7.10.7), an appropriate segment address must be explicitly provided. The **reti** instruction has operands similar to the **ret** instruction.

7.8.6 Input/Output Instructions

The XENIX assembler has modified the **in** and **out** instructions to include the new forms **inw** and **outw**. The **in** and **out** instructions now operate strictly on byte values; **inw** and **outw** operate on words. Furthermore, these instructions take only one operand – the port number. The **al** or **ax** register is accessed as appropriate and must not be given as an operand.

7.8.7 80286 Instructions

Assembly language programmers who wish to use 80286 instructions can insert

the binary opcode of a given instruction into the instruction stream using the **.byte** and **.word** directives. For example, the directive

places the binary opcode of the **pushi** (for "push"immediate instruction in the instruction stream. This code is equivalent to the 80286 instruction

```
pushi *1
```

A programmer can also use the power of the C language preprocessor to create macros for the 80286 instructions. For example, if the macro definition

```
#define PUSHIB(x)          .byte /6a, *x
```

is included in an assembly language source, then the macro call

```
PUSHIB(1)
```

may be used for a **pushi** instruction in place of the **.byte** directive. In this case, you must invoke the C proprocessor using the cc command to resolve this macro.

Note that priviledged 80286 instructions are not available to user programs.

7.9 Addressing Modes

The XENIX 8086/286 assembler provides many different ways to access instruction operands. Operands may be contained in registers, within the instruction itself, in memory, or in I/O ports. In addition, the addresses of memory and I/O port operands can be calculated in several different ways.

The following sections describe the format and meaning of instruction operands.

7.9.1 Register Operands

Register operands are the 8086/286 CPU registers. In an instruction, a register operand causes the instruction's action to be performed on the contents of the register. A register operand may be any one of the following:

ax	ah	al
bx	bh	bl
cx	ch	cl
dx	dh	dl
di	si	bp
sp	cs	ds
ss	es	

Register operands may be used for the source or destination in an instruction.

Since these operands are encoded in a few bits, instructions that specify only register operands are generally the most compact. They are also the fastest, since operations on registers are performed entirely within the CPU.

The following are examples of register operands.

```
sub     ax,bx
addb    ah,dl
cmp     cs,ds
```

7.9.2 Immediate Operands

Immediate operands are byte or word constants given with the instruction itself. These operands have the forms

```
* expression
# expression
```

where * specifies a byte constant, # specifies a word constant, and *expression* is an absolute expression or a symbol which defines the constant's value.

Since immediate operands are constants, they cannot be used as the destination operand. Note that the assembler does not check the operand size.

The following examples illustrate immediate operands.

```
movb    cx,*33
mov     cx,#(122/2)
addb    ax,*NAME
neg     TOP
```

7.9.3 Direct Address Operands

Direct address operands are the bytes or words in memory at the given direct addresses. Direct addresses have the form

```
expression
```

where *expression* is an absolute expression or symbol which resolves to a memory address. The direct address gives the location of the operand in terms of an offset from the beginning of the current segment or the segment in which the given symbol is defined.

Direct address operands may be used for source or destination. Although absolute addresses are allowed, symbols should be used whenever possible. The size of the operand depends on the instruction in which it is used.

The following examples illustrate direct address operands.

```
mov     cx,free
movb    Darray+204,*1
```

7.9.4 Based Operands

Based operands are bytes or words in memory whose address is computed by adding a constant and one of the base registers **bp** or **bx**. The operands have the forms

[*expression*] (bp)
[*expression*] (bx)

where *expression* is an absolute expression or symbol which resolves to an absolute.

Based operands are typically used to access structures. The base register points to the start of the structure and items in the structure are addressed by the giving an appropriate *expression*.

The following examples illustrate based operands.

```
mov     2(bp), #1000
movb    ax, TOP(bx)
neg     -4(bp)
mov     ax, (bx)(bp)
```

7.9.5 Indexed Operands

Indexed operand are bytes or words in memory whose addresses are computed by adding a constant and one of the index registers **di** or **si**. The operands have the forms

[*expression*] (di)
[*expression*] (si)

where *expression* is an absolute expression or symbol which resolves to an absolute.

Indexed operands are often used to access elements in an array. The *expression* points to the start of the array. The index register is given the index value of the element to be accessed. Since all array elements are the same length, simple arithmetic on the index register will select any element.

The following examples illustrate indexed operands.

```
movb    ax,Darray(di)
addb    4096(si),*1
```

7.9.6 Based Indexed Operands

Based indexed operands are bytes or words in memory whose addresses are computed by adding a constant, a base register, and an index register. The operands have the forms

```
[ expression ] (bx) (di)
[ expression ] (bx) (si)
[ expression ] (bp) (di)
[ expression ] (bp) (si)
```

where *expression* is an absolute expression or a symbol which resolves to an absolute.

Based indexed operands provide a very flexible method of accessing items that require two address components. For example, elements of multidimensional arrays can be accessed by setting *expression* to the start address of the array and assigning the appropriately scaled index values to the base and index registers.

The following examples illustrate based indexed operands.

```
movb    Darray(bx)(di),*1
mov     ax, (bx)(si)
neg     -2(bp)(si)
```

7.9.7 Indirect Address Operands

Indirect address operands are instruction addresses that are stored in memory at given indirect addresses. Indirect address operands have the form

```
@expression
```

where *expression* is an absolute expression or a symbol that resolves to an absolute.

Indirect address operands may be used only with the **calli**, **call**, **jmpi**, and **jmp** instructions. When used with the intersegment call or jump instruction, *expression* must point to a 4-byte segment/offset instruction address. When used with the call or jump instruction, *expression* must point to a 2-byte offset to an instruction.

The following examples illustrate indirect address operands.

```
moncall:
        .word   OFFSET
        .word   SEG_SELECT
subtest:
        .word   OFFSET
        .text
examp:
        calli   @moncall
        jmp     @subtest
```

7.10 Diagnostics

When syntactic errors occur, the assembler displays the line number and the name of the file containing the error. If the errors are encountered in the first pass of the assembler, the second pass is canceled and no object file is created.

Error messages have the following form

ERROR syntax error, line *nnn*
file: *eee* errors

where *nnn* is the line number(s) containing the error, *file* is the name of the file, and *eee* is the total number of errors.

Chapter 8
Lex: A Lexical Analyzer

8.1 Introduction 8-1

8.2 Lex Source Format 8-2

8.3 Lex Regular Expressions 8-3

8.4 Invoking *lex* 8-4

8.5 Specifying Character Classes 8-5

8.6 Specifying an Arbitrary Character 8-6

8.7 Specifying Optional Expressions 8-6

8.8 Specifying Repeated Expressions 8-6

8.9 Specifying Alternation and Grouping 8-7

8.10 Specifying Context Sensitivity 8-7

8.11 Specifying Expression Repetition 8-8

8.12 Specifying Definitions 8-8

8.13 Specifying Actions 8-8

8.14 Handling Ambiguous Source Rules 8-12

8.15 Specifying Left Context Sensitivity 8-15

8.16 Specifying Source Definitions 8-17

8.17 Lex and Yacc 8-18

8.18 Specifying Character Sets 8-22

8.19 Source Format 8-23

8.1 Introduction

Lex is a program generator designed for lexical processing of character input streams. It accepts a high-level, problem-oriented specification for character string matching, and produces a C program that recognizes regular expressions. The regular expressions are specified by the user in the source specifications given to lex. The lex code recognizes these expressions in an input stream and partitions the input stream into strings matching the expressions. At the boundaries between strings, program sections provided by the user are executed. The lex source file associates the regular expressions and the program fragments. As each expression appears in the input to the program written by lex, the corresponding fragment is executed.

The user supplies the additional code needed to complete his tasks, including code written by other generators. The program that recognizes the expressions is generated in the from the user's C program fragments. Lex is not a complete language, but rather a generator representing a new language feature added on top of the C programming language.

Lex turns the user's expressions and actions (called *source* in this chapter) into a C program named *yylex*. The *yylex* program recognizes expressions in a stream (called input in this chapter) and performs the specified actions for each expression as it is detected.

Consider a program to delete from the input all blanks or tabs at the ends of lines. The following lines

```
%%
[\t]+$   ;
```

are all that is required. The program contains a %% delimiter to mark the beginning of the rules, and one rule. This rule contains a regular expression that matches one or more instances of the characters blank or tab (written \t for visibility, in accordance with the C language convention) just prior to the end of a line. The brackets indicate the character class made of blank and tab; the + indicates one or more of the previous item; and the dollar sign ($) indicates the end of the line. No action is specified, so the program generated by lex will ignore these characters. Everything else will be copied. To change any remaining string of blanks or tabs to a single blank, add another rule:

```
%%
[\t]+$   ;
[\t]+    printf(" ");
```

The finite automaton generated for this source scans for both rules at once, observes at the termination of the string of blanks or tabs whether or not there is a newline character, and then executes the desired rule's action. The first rule matches all strings of blanks or tabs at the end of lines, and the second rule matches all remaining strings of blanks or tabs.

Lex can be used alone for simple transformations, or for analysis and statistics gathering on a lexical level. Lex can also be used with a parser generator to perform the lexical analysis phase; it is especially easy to interface lex and yacc. Lex programs recognize only regular expressions; yacc writes parsers that accept a large class of context-free grammars, but that require a lower level analyzer to recognize input tokens. Thus, a combination of lex and yacc is often appropriate. When used as a preprocessor for a later parser generator, lex is used to partition the input stream, and the parser generator assigns structure to the resulting pieces. Additional programs, written by other generators or by hand, can be added easily to programs written by lex. Yacc users will realize that the name *yylex* is what yacc expects its lexical analyzer to be named, so that the use of this name by lex simplifies interfacing.

Lex generates a deterministic finite automaton from the regular expressions in the source. The automaton is interpreted, rather than compiled, in order to save space. The result is still a fast analyzer. In particular, the time taken by a lex program to recognize and partition an input stream is proportional to the length of the input. The number of lex rules or the complexity of the rules is not important in determining speed, unless rules which include forward context require a significant amount of rescanning. What does increase with the number and complexity of rules is the size of the finite automaton, and therefore the size of the program generated by lex.

In the program written by lex, the user's fragments (representing the actions to be performed as each regular expression is found) are gathered as cases of a switch. The automaton interpreter directs the control flow. Opportunity is provided for the user to insert either declarations or additional statements in the routine containing the actions, or to add subroutines outside this action routine.

Lex is not limited to source that can be interpreted on the basis of one character lookahead. For example, if there are two rules, one looking for *ab* and another for *abcdefg*, and the input stream is *abcdefh*, lex will recognize *ab* and leave the input pointer just before *cd*. Such backup is more costly than the processing of simpler languages.

8.2 Lex Source Format

The general format of lex source is:

{definitions}
%%
{rules}
%%
{user subroutines}

where the definitions and the user subroutines are often omitted. The second %% is optional, but the first is required to mark the beginning of the rules. The absolute minimum lex program is thus

```
%%
```

(no definitions, no rules) which translates into a program that copies the input to the output unchanged.

In the lex program format shown above, the rules represent the user's control decisions. They make up a table in which the left column contains regular expressions and the right column contains actions, program fragments to be executed when the expressions are recognized. Thus the following individual rule might appear:

```
integer    printf("found keyword INT");
```

This looks for the string *integer* in the input stream and prints the message

```
found keyword INT
```

whenever it appears in the input text. In this example the C library function *printf*() is used to print the string. The end of the lex regular expression is indicated by the first blank or tab character. If the action is merely a single C expression, it can be given on the right side of the line; if it is compound, or takes more than a line, it should be enclosed in braces. As a slightly more useful example, suppose it is desired to change a number of words from British to American spelling. Lex rules such as

```
colour          printf("color");
mechanise       printf("mechanize");
petrol          printf("gas");
```

would be a start. These rules are not quite enough, since the word *petroleum* would become *gaseum*; a way of dealing with such problems is described in a later section.

8.3 Lex Regular Expressions

A regular expression specifies a set of strings to be matched. It contains text characters (that match the corresponding characters in the strings being compared) and operator characters (these specify repetitions, choices, and other features). The letters of the alphabet and the digits are always text characters. Thus, the regular expression

```
integer
```

matches the string *integer* wherever it appears and the expression

```
a57D
```

looks for the string *a57D*.

The operator characters are

" \ [ ] ^ - ? . * + | ( ) $ / { } % < >

If any of these characters are to be used literally, they needed to be quoted individually with a backslash (\) or as a group within quotation marks ("). The quotation mark operator (") indicates that whatever is contained between a pair of quotation marks is to be taken as text characters. Thus

```
xyz"++"
```

matches the string *xyz++* when it appears. Note that a part of a string may be quoted. It is harmless but unnecessary to quote an ordinary text character; the expression

```
"xyz++"
```

is the same as the one above. Thus by quoting every nonalphanumeric character being used as a text character, you need not memorize the above list of current operator characters.

An operator character may also be turned into a text character by preceding it with a backslash (\) as in

```
xyz\+\+
```

which is another, less readable, equivalent of the above expressions. The quoting mechanism can also be used to get a blank into an expression; normally, as explained above, blanks or tabs end a rule. Any blank character not contained within brackets must be quoted. Several normal C escapes with the backslash (\ ) are recognized:

\n newline

\t tab

\b backspace

\\ backslash

Since newline is illegal in an expression, a *n* must be used; it is not required to escape tab and backspace. Every character but blank, tab, newline and the list above is always a text character.

8.4 Invoking *lex*

There are two steps in compiling a lex source program. First, the lex source must be turned into a generated program in the host general purpose language. Then this program must be compiled and loaded, usually with a library of lex

subroutines. The generated program is in a file named lex.yy.c. The I/O library is defined in terms of the C standard library.

The library is accessed by the loader flag *–ll*. So an appropriate set of commands is

```
lex source
cc lex.yy.c –ll
```

The resulting program is placed on the usual file *a.out* for later execution. To use lex with yacc see the section "Lex and Yacc" in this chapter and Chapter 9, "Yacc: A Compiler-Compiler"". Although the default lex I/O routines use the C standard library, the lex automata themselves do not do so. If private versions of *input*, *output*, and *unput* are given, the library can be avoided.

8.5 Specifying Character Classes

Classes of characters can be specified using brackets: [and]. The construction

```
[abc]
```

matches a single character, which may be *a*, *b*, or *c*. Within square brackets, most operator meanings are ignored. Only three characters are special: these are the backslash (\), the dash (-), and the caret (^). The dash character indicates ranges. For example

```
[a-z0-9<>_]
```

indicates the character class containing all the lowercase letters, the digits, the angle brackets, and underline. Ranges may be given in either order. Using the dash between any pair of characters that are not both uppercase letters, both lowercase letters, or both digits is implementation dependent and causes a warning message. If it is desired to include the dash in a character class, it should be first or last; thus

```
[-+0-9]
```

matches all the digits and the plus and minus signs.

In character classes, the caret (^) operator must appear as the first character after the left bracket; it indicates that the resulting string is to be complemented with respect to the computer character set. Thus

```
[^abc]
```

matches all characters except *a*, *b*, or *c*, including all special or control characters; or

```
[^a-zA-Z]
```

is any character which is not a letter. The backslash (\) provides an escape mechanism within character class brackets, so that characters can be entered literally by preceding them with this character.

8.6 Specifying an Arbitrary Character

To match almost any character, the period (.) designates the class of all characters except a newline. Escaping into octal is possible although nonportable. For example

```
[\40-\176]
```

matches all printable characters in the ASCII character set, from octal 40 (blank) to octal 176 (tilde).

8.7 Specifying Optional Expressions

The question mark (?) operator indicates an optional element of an expression. Thus

```
ab?c
```

matches either *ac* or *abc*. Note that the meaning of the question mark here differs from its meaning in the shell.

8.8 Specifying Repeated Expressions

Repetitions of classes are indicated by the asterisk (*) and plus (+) operators. For example

```
a*
```

matches any number of consecutive *a* characters, including zero; while *a+* matches one or more instances of *a*. For example,

```
[a-z]+
```

matches all strings of lowercase letters, and

```
[A-Za-z][A-Za-z0-9]*
```

matches all alphanumeric strings with a leading alphabetic character; this is a typical expression for recognizing identifiers in computer languages.

8.9 Specifying Alternation and Grouping

The vertical bar (|) operator indicates alternation. For example

```
(ab|cd)
```

matches either *ab* or *cd*. Note that parentheses are used for grouping, although they are not necessary at the outside level. For example

```
ab|cd
```

would have sufficed in the preceding example. Parentheses should be used for more complex expressions, such as

```
(ab|cd+)?(ef)*
```

which matches such strings as *abefef*, *efefef*, *cdef*, and *cddd*, but not *abc*, *abcd*, or *abcdef*.

8.10 Specifying Context Sensitivity

Lex recognizes a small amount of surrounding context. The two simplest operators for this are the caret (^) and the dollar sign (**$**). If the first character of an expression is a caret, then the expression is only matched at the beginning of a line (after a newline character, or at the beginning of the input stream). This can never conflict with the other meaning of the caret, complementation of character classes, since complementation only applies within brackets. If the very last character is dollar sign, the expression only matched at the end of a line (when immediately followed by newline). The latter operator is a special case of the slash (/) operator, which indicates trailing context. The expression

```
ab/cd
```

matches the string *ab*, but only if followed by *cd*. Thus

```
ab$
```

is the same as

```
ab/\n
```

Left context is handled in lex by specifying start conditions as explained in the section "Specifying Left Context Sensitivity". If a rule is only to be executed when the lex automaton interpreter is in start condition *x*, the rule should be enclosed in angle brackets:

```
<x>
```

If we considered being at the beginning of a line to be start condition ONE, then the caret (^) operator would be equivalent to

```
<ONE>
```

Start conditions are explained more fully later.

8.11 Specifying Expression Repetition

The curly braces ({ and }) specify either repetitions (if they enclose numbers) or definition expansion (if they enclose a name). For example

```
{digit}
```

looks for a predefined string named *digit* and inserts it at that point in the expression.

8.12 Specifying Definitions

The definitions are given in the first part of the lex input, before the rules. In contrast,

```
a{1,5}
```

looks for 1 to 5 occurrences of the character *a*.

Finally, an initial percent sign (%) is special, since it is the separator for lex source segments.

8.13 Specifying Actions

When an expression is matched by a pattern of text in the input, lex executes the corresponding action. This section describes some features of lex which aid in writing actions. Note that there is a default action, which consists of copying the input to the output. This is performed on all strings not otherwise matched. Thus the lex user who wishes to absorb the entire input, without producing any output, must provide rules to match everything. When lex is being used with yacc, this is the normal situation. You may consider that actions are what is done instead of copying the input to the output; thus, in general, a rule which merely copies can be omitted.

One of the simplest things that can be done is to ignore the input. Specifying a C null statement ; as an action causes this result. A frequent rule is

```
[ \t\n]   ;
```

which causes the three spacing characters (blank, tab, and newline) to be

ignored.

Another easy way to avoid writing actions is to use the repeat action character, |, which indicates that the action for this rule is the action for the next rule. The previous example could also have been written

```
" "          |
"\t"         |
"\n"         ;
```

with the same result, although in a different style. The quotes around \n and \t are not required.

In more complex actions, you often want to know the actual text that matched some expression like:

```
[a-z]+
```

Lex leaves this text in an external character array named *yytext*. Thus, to print the name found, a rule like

```
[a-z]+    printf("%s", yytext);
```

prints the string in *yytext*. The C function *printf* accepts a format argument and data to be printed; in this case, the format is *print string* where the percent sign (%) indicates data conversion, and the *s* indicate string type, and the data are the characters in *yytext*. So this just places the matched string on the output. This action is so common that it may be written as ECHO. For example

```
[a-z]+    ECHO;
```

is the same as the preceding example. Since the default action is just to print the characters found, one might ask why give a rule, like this one, which merely specifies the default action? Such rules are often required to avoid matching some other rule that is not desired. For example, if there is a rule that matches *read* it will normally match the instances of *read* contained in *bread* or *readjust*; to avoid this, a rule of the form

```
[a-z]+
```

is needed. This is explained further below.

Sometimes it is more convenient to know the end of what has been found; hence lex also provides a count of the number of characters matched in the variable, *yyleng*. To count both the number of words and the number of characters in words in the input, you might write

```
[a-zA-Z]+        {words++; chars += yyleng;}
```

which accumulates in the variables *chars* the number of characters in the words

recognized. The last character in the string matched can be accessed with:

```
yytext[yyleng-1]
```

Occasionally, a lex action may decide that a rule has not recognized the correct span of characters. Two routines are provided to aid with this situation. First, *yymore*() can be called to indicate that the next input expression recognized is to be tacked on to the end of this input. Normally, the next input string will overwrite the current entry in *yytext*. Second, *yyless*(n) may be called to indicate that not all the characters matched by the currently successful expression are wanted right now. The argument *n* indicates the number of characters in *yytext* to be retained. Further characters previously matched are returned to the input. This provides the same sort of lookahead offered by the slash (/) operator, but in a different form.

For example, consider a language that defines a string as a set of characters between quotation marks ("), and provides that to include a quotation mark in a string, it must be preceded by a backslash (\). The regular expression that matches this is somewhat confusing, so that it might be preferable to write

```
\"[^"]*  {
         if (yytext[yyleng-1] == '\\')
             yymore();
         else
             ... normal user processing
         }
```

which, when faced with a string such as

```
"abc\"def"
```

will first match the five characters

```
"abc\
```

and then the call to *yymore*() will cause the next part of the string,

```
"def
```

to be tacked on the end. Note that the final quotation mark terminating the string should be picked up in the code labeled normal processing.

The function *yyless*() might be used to reprocess text in various circumstances. Consider the problem in the older C syntax of distinguishing the ambiguity of =-*a*. Suppose it is desired to treat this as =- *a* and to print a message. A rule might be

```
=-[a-zA-Z]        {
        printf("Operator (=-) ambiguous\n");
        yyless(yyleng-1);
        ... action for =- ...
        }
```

which prints a message, returns the letter after the operator to the input stream, and treats the operator as =-.

Alternatively it might be desired to treat this as = -*a*. To do this, just return the minus sign as well as the letter to the input. The following performs the interpretation:

```
=-[a-zA-Z]        {
        printf("Operator (=-) ambiguous\n");
        yyless(yyleng-2);
        ... action for = ...
        }
```

Note that the expressions for the two cases might more easily be written

```
=-/[A-Za-z]
```

in the first case and

```
=/-[A-Za-z]
```

in the second: no backup would be required in the rule action. It is not necessary to recognize the whole identifier to observe the ambiguity. The possibility of =-*3*, however, makes

```
=-/[^ \t\n]
```

a still better rule.

In addition to these routines, lex also permits access to the I/O routines it uses. They include:

1. *input*() which returns the next input character;
2. *output*(c) which writes the character *c* on the output; and
3. *unput*(c) which pushes the character *c* back onto the input stream to be read later by *input*().

By default these routines are provided as macro definitions, but the user can override them and supply private versions. These routines define the relationship between external files and internal characters, and must all be retained or modified consistently. They may be redefined, to cause input or

output to be transmitted to or from strange places, including other programs or internal memory; but the character set used must be consistent in all routines; a value of zero returned by *input* must mean end-of-file; and the relationship between *unput* and *input* must be retained or the lookahead will not work. Lex does not look ahead at all if it does not have to, but every rule containing a slash (/) or ending in one of the following characters implies lookahead:

```
+ * ? $
```

Lookahead is also necessary to match an expression that is a prefix of another expression. See below for a discussion of the character set used by lex. The standard lex library imposes a 100 character limit on backup.

Another lex library routine that you sometimes want to redefine is *yywrap*() which is called whenever lex reaches an end-of-file. If *yywrap* returns a 1, lex continues with the normal wrapup on end of input. Sometimes, however, it is convenient to arrange for more input to arrive from a new source. In this case, the user should provide a *yywrap* that arranges for new input and returns 0. This instructs lex to continue processing. The default *yywrap* always returns 1.

This routine is also a convenient place to print tables, summaries, etc. at the end of a program. Note that it is not possible to write a normal rule that recognizes end-of-file; the only access to this condition is through *yywrap*(). In fact, unless a private version of *input*() is supplied a file containing nulls cannot be handled, since a value of 0 returned by *input* is taken to be end-of-file.

8.14 Handling Ambiguous Source Rules

Lex can handle ambiguous specifications. When more than one expression can match the current input, lex chooses as follows:

- The longest match is preferred.
- Among rules that match the same number of characters, the first given rule is preferred.

For example, suppose the following rules are given:

```
integer    keyword action ...;
[a-z]+     identifier action ...;
```

If the input is *integers*, it is taken as an identifier, because

```
[a-z]+
```

matches 8 characters while

```
integer
```

matches only 7. If the input is *integer*, both rules match 7 characters, and the keyword rule is selected because it was given first. Anything shorter (e.g., *int*) does not match the expression *integer*, so the identifier interpretation is used.

The principle of preferring the longest match makes certain constructions dangerous, such as the following:

```
.*
```

For example

```
'.*'
```

might seem a good way of recognizing a string in single quotes. But it is an invitation for the program to read far ahead, looking for a distant single quote. Presented with the input

```
'first' quoted string here, 'second' here
```

the above expression matches

```
'first' quoted string here, 'second'
```

which is probably not what was wanted. A better rule is of the form

```
'[^'\n]*'
```

which, on the above input, stops after 'first'. The consequences of errors like this are mitigated by the fact that the dot (.) operator does not match a newline. Therefore, no more than one line is ever matched by such expressions. Don't try to defeat this with expressions like

```
[.\n]+
```

or their equivalents: the lex generated program will try to read the entire input file, causing internal buffer overflows.

Note that lex is normally partitioning the input stream, not searching for all possible matches of each expression. This means that each character is accounted for once and only once. For example, suppose it is desired to count occurrences of both *she* and *he* in an input text. Some lex rules to do this might be

```
she     s++;
he      h++;
\n      |
.       ;
```

where the last two rules ignore everything besides *he* and *she*. Remember that the period (.) does not include the newline. Since *she* includes *he*, lex will normally not recognize the instances of *he* included in *she*, since once it has passed a *she* those characters are gone.

Sometimes the user would like to override this choice. The action REJECT means go do the next alternative. It causes whatever rule was second choice after the current rule to be executed. The position of the input pointer is adjusted accordingly. Suppose the user really wants to count the included instances of *he*:

```
she     {s++; REJECT;}
he      {h++; REJECT;}
\n      |
.       ;
```

These rules are one way of changing the previous example to do just that. After counting each expression, it is rejected; whenever appropriate, the other expression will then be counted. In this example, of course, the user could note that *she* includes *he*, but not vice versa, and omit the REJECT action on *he*; in other cases, however, it would not be possible to tell which input characters were in both classes.

Consider the two rules

```
a[bc]+   { ... ; REJECT;}
a[cd]+   { ... ; REJECT;}
```

If the input is *ab*, only the first rule matches, and on *ad* only the second matches. The input string *accb* matches the first rule for four characters and then the second rule for three characters. In contrast, the input *accd* agrees with the second rule for four characters and then the first rule for three.

In general, REJECT is useful whenever the purpose of lex is not to partition the input stream but to detect all examples of some items in the input, and the instances of these items may overlap or include each other. Suppose a digram table of the input is desired; normally the digrams overlap, that is the word *the* is considered to contain both *th* and *he*. Assuming a two-dimensional array named *digram* to be incremented, the appropriate source is

```
%%
[a-z][a-z] {digram[yytext[0]][yytext[1]]++; REJECT;}
.       ;
\n      ;
```

where the REJECT is necessary to pick up a letter pair beginning at every character, rather than at every other character.

Remember that REJECT does not rescan the input. Instead it remembers the results of the previous scan. This means that if a rule with trailing context is

found, and REJECT executed, you must not have used *unput* to change the characters forthcoming from the input stream. This is the only restriction to ability to manipulate the not-yet-processed input.

8.15 Specifying Left Context Sensitivity

Sometimes it is desirable to have several sets of lexical rules to be applied at different times in the input. For example, a compiler preprocessor might distinguish preprocessor statements and analyze them differently from ordinary statements. This requires sensitivity to prior context, and there are several ways of handling such problems. The caret (^) operator, for example, is a prior context operator, recognizing immediately preceding left context just as the dollar sign ($) recognizes immediately following right context. Adjacent left context could be extended, to produce a facility similar to that for adjacent right context, but it is unlikely to be as useful, since often the relevant left context appeared some time earlier, such as at the beginning of a line.

This section describes three means of dealing with different environments:

1. The use of flags, when only a few rules change from one environment to another
2. The use of start conditions with rules
3. The use multiple lexical analyzers running together.

In each case, there are rules that recognize the need to change the environment in which the following input text is analyzed, and set some parameter to reflect the change. This may be a flag explicitly tested by the user's action code; such a flag is the simplest way of dealing with the problem, since lex is not involved at all. It may be more convenient, however, to have lex remember the flags as initial conditions on the rules. Any rule may be associated with a start condition. It will only be recognized when lex is in that start condition. The current start condition may be changed at any time. Finally, if the sets of rules for the different environments are very dissimilar, clarity may be best achieved by writing several distinct lexical analyzers, and switching from one to another as desired.

Consider the following problem: copy the input to the output, changing the word *magic* to *first* on every line that began with the letter *a*, changing *magic* to *second* on every line that began with the letter *b*, and changing *magic* to *third* on every line that began with the letter *c*. All other words and all other lines are left unchanged.

These rules are so simple that the easiest way to do this job is with a flag:

```
            int flag;
%%
^a          {flag = 'a'; ECHO;}
^b          {flag = 'b'; ECHO;}
^c          {flag = 'c'; ECHO;}
\n          {flag = 0 ; ECHO;}
magic       {
            switch (flag)
            {
            case 'a': printf("first"); break;
            case 'b': printf("second"); break;
            case 'c': printf("third"); break;
            default: ECHO; break;
            }
            }
```

should be adequate.

To handle the same problem with start conditions, each start condition must be introduced to lex in the definitions section with a line reading

%Start name1 name2 ...

where the conditions may be named in any order. The word *Start* may be abbreviated to *s* or *S*. The conditions may be referenced at the head of a rule with angle brackets. For example

<name1>expression

is a rule that is only recognized when lex is in the start condition *name1*. To enter a start condition, execute the action statement

BEGIN name1;

which changes the start condition to *name1*. To return to the initial state

BEGIN 0;

resets the initial condition of the lex automaton interpreter. A rule may be active in several start conditions; for example:

<name1,name2,name3>

is a legal prefix. Any rule not beginning with the <> prefix operator is always active.

The same example as before can be written:

```
%START AA BB CC
%%
^a      {ECHO; BEGIN AA;}
^b      {ECHO; BEGIN BB;}
^c      {ECHO; BEGIN CC;}
\n      {ECHO; BEGIN 0;}
<AA>magic       printf("first");
<BB>magic       printf("second");
<CC>magic       printf("third");
```

where the logic is exactly the same as in the previous method of handling the problem, but lex does the work rather than the user's code.

8.16 Specifying Source Definitions

Remember the format of the lex source:

```
{definitions}
%%
{rules}
%%
{user routines}
```

So far only the rules have been described. You will need additional options, though, to define variables for use in your program and for use by lex. These can go either in the definitions section or in the rules section.

Remember that lex is turning the rules into a program. Any source not intercepted by lex is copied into the generated program. There are three classes of such things:

1. Any line that is not part of a lex rule or action which begins with a blank or tab is copied into the lex generated program. Such source input prior to the first %% delimiter will be external to any function in the code; if it appears immediately after the first %%, it appears in an appropriate place for declarations in the function written by lex which contains the actions. This material must look like program fragments, and should precede the first lex rule.

   As a side effect of the above, lines that begin with a blank or tab, and which contain a comment, are passed through to the generated program. This can be used to include comments in either the lex source or the generated code. The comments should follow the conventions of the C language.

2. Anything included between lines containing only %{ and %} is copied out as above. The delimiters are discarded. This format permits entering text like preprocessor statements that must begin in column

1, or copying lines that do not look like programs.

3. Anything after the third %% delimiter, regardless of formats, is copied out after the lex output.

Definitions intended for lex are given before the first %% delimiter. Any line in this section not contained between %{ and %}, and beginning in column 1, is assumed to define lex substitution strings. The format of such lines is

```
name translation
```

and it causes the string given as a translation to be associated with the name. The name and translation must be separated by at least one blank or tab, and the name must begin with a letter. The translation can then be called out by the {name} syntax in a rule. Using {D} for the digits and {E} for an exponent field, for example, might abbreviate rules to recognize numbers:

```
D                       [0-9]
E                       [DEde][-+]?{D}+
%%
{D}+                    printf("integer");
{D}+"."{D}*({E})?       |
{D}*"."{D}+({E})?       |
{D}+{E}                         printf("real");
```

Note the first two rules for real numbers; both require a decimal point and contain an optional exponent field, but the first requires at least one digit before the decimal point and the second requires at least one digit after the decimal point. To correctly handle the problem posed by a FORTRAN expression such as *35.EQ.I*, which does not contain a real number, a context-sensitive rule such as

```
[0-9]+/"."EQ      printf("integer");
```

could be used in addition to the normal rule for integers.

The definitions section may also contain other commands, including a character set table, a list of start conditions, or adjustments to the default size of arrays within lex itself for larger source programs. These possibilities are discussed in the section "Source Format".

8.17 Lex and Yacc

If you want to use lex with yacc, note that what lex writes is a program named *yylex*(), the name required by yacc for its analyzer. Normally, the default main program on the lex library calls this routine, but if yacc is loaded, and its main program is used, yacc will call *yylex*(). In this case, each lex rule should end with

```
return(token);
```

where the appropriate token value is returned. An easy way to get access to **yacc**'s names for tokens is to compile the **lex** output file as part of the **yacc** output file by placing the line

```
# include "lex.yy.c"
```

in the last section of **yacc** input. Supposing the grammar to be named *good* and the lexical rules to be named *better* the XENIX command sequence can just be:

```
yacc good
lex better
cc y.tab.c -ly -ll
```

The **yacc** library (–ly) should be loaded before the **lex** library, to obtain a main program which invokes the **yacc** parser. The generation of **lex** and **yacc** programs can be done in either order.

As a trivial problem, consider copying an input file while adding 3 to every positive number divisible by 7. Here is a suitable **lex** source program to do just that:

```
%%
        int k;
[0-9]+  {
        k = atoi(yytext);
        if (k%7 == 0)
            printf("%d", k+3);
        else
            printf("%d",k);
        }
```

The rule [0–9]+ recognizes strings of digits; *atoi*() converts the digits to binary and stores the result in *k*. The remainder operator (%) is used to check whether *k* is divisible by 7; if it is, it is incremented by 3 as it is written out. It may be objected that this program will alter such input items as 49.63 or X7. Furthermore, it increments the absolute value of all negative numbers divisible by 7. To avoid this, just add a few more rules after the active one, as here:

```
%%
        int k;
-?[0-9]+  {
        k = atoi(yytext);
        printf("%d", k%7 == 0 ? k+3 : k);
        }
-?[0-9.]+       ECHO;
[A-Za-z][A-Za-z0-9]+    ECHO;
```

Numerical strings containing a decimal point or preceded by a letter will be

picked up by one of the last two rules, and not changed. The **if–else** has been replaced by a C conditional expression to save space; the form *a?b:c* means: if *a* then *b* else *c*.

For an example of statistics gathering, here is a program which makes histograms of word lengths, where a word is defined as a string of letters.

```
        int lengs[100];
%%
[a-z]+  lengs[yyleng]++;
.       |
\n      ;
%%
yywrap()
{
int i;
printf("Length  No. words\n");
for(i=0; i<100; i++)
    if (lengs[i] > 0)
        printf("%5d%10d\n",i,lengs[i]);
return(1);
}
```

This program accumulates the histogram, while producing no output. At the end of the input it prints the table. The final statement *return*(1); indicates that lex is to perform wrapup. If *yywrap*() returns zero (false) it implies that further input is available and the program is to continue reading and processing. To provide a *yywrap*() that never returns true causes an infinite loop.

As a larger example, here are some parts of a program written to convert double precision FORTRAN to single precision FORTRAN. Because FORTRAN does not distinguish between upper- and lowercase letters, this routine begins by defining a set of classes including both cases of each letter:

```
a       [aA]
b       [bB]
c       [cC]
.       .
.       .
.       .
z       [zZ]
```

An additional class recognizes white space:

```
W       [ \t]*
```

The first rule changes *double precision* to *real*, or *DOUBLE PRECISION* to *REAL*.

```
{d}{o}{u}{b}{l}{e}{W}{p}{r}{e}{c}{i}{s}{i}{o}{n} {
    printf(yytext[0]=='d'? "real" : "REAL");
    }
```

Care is taken throughout this program to preserve the case of the original program. The conditional operator is used to select the proper form of the keyword. The next rule copies continuation card indications to avoid confusing them with constants:

```
^"     "[^ 0]        ECHO;
```

In the regular expression, the quotes surround the blanks. It is interpreted as beginning of line, then five blanks, then anything but blank or zero." Note the two different meanings of the caret (^) here. There follow some rules to change double precision constants to ordinary floating constants.

```
[0-9]+{W}{d}{W}[+-]?{W}[0-9]+          |
[0-9]+{W}"."{W}{d}{W}[+-]?{W}[0-9]+    |
"."{W}[0-9]+{W}{d}{W}[+-]?{W}[0-9]+    {
    /* convert constants */
    for(p=yytext; *p != 0; p++)
        {
        if (*p == 'd' || *p == 'D')
            *p+= 'e'- 'd';
        ECHO;
        }
```

After the floating point constant is recognized, it is scanned by the **for** loop to find the letter "d" or "D". The program then adds "'e' –' d' " which converts it to the next letter of the alphabet. The modified constant, now single precision, is written out again. There follow a series of names which must be respelled to remove their initial "d". By using the array *yytext* the same action suffices for all the names (only a sample of a rather long list is given here).

```
{d}{s}{i}{n}      |
{d}{c}{o}{s}      |
{d}{s}{q}{r}{t}   |
{d}{a}{t}{a}{n}   |
...
{d}{f}{l}{o}{a}{t}          printf("%s",yytext+1);
```

Another list of names must have initial *d* changed to initial *a*:

```
{d}{l}{o}{g}      |
{d}{l}{o}{g}10    |
{d}{m}{i}{n}1     |
{d}{m}{a}{x}1     {
        yytext[0] += 'a' - 'd';
        ECHO;
        }
```

And one routine must have initial *d* changed to initial *r*:

```
{d}1{m}{a}{c}{h}        {
        yytext[0] += 'r' - 'd';
        ECHO;
}
```

To avoid such names as *dsinx* being detected as instances of *dsin*, some final rules pick up longer words as identifiers and copy some surviving characters:

```
[A-Za-z][A-Za-z0-9]*    |
[0-9]+                  |
\n                      |
.       ECHO;
```

Note that this program is not complete; it does not deal with the spacing problems in FORTRAN or with the use of keywords as identifiers.

8.18 Specifying Character Sets

The programs generated by lex handle character I/O only through the routines *input*, *output*, and *unput*. Thus the character representation provided in these routines is accepted by lex and employed to return values in *yytext*. For internal use a character is represented as a small integer which, if the standard library is used, has a value equal to the integer value of the bit pattern representing the character on the host computer. Normally, the letter *a* is represented as the same form as the character constant:

```
'a'
```

If this interpretation is changed, by providing I/O routines which translate the characters, lex must be told about it, by giving a translation table. This table must be in the definitions section, and must be bracketed by lines containing only *%T*. The table contains lines of the form

```
{integer} {character string}
```

which indicate the value associated with each character. For example:

```
%T
 1        Aa
 2        Bb
...
26        Zz
27        \n
28        +
29        -
30        0
31        1
...
39        9
%T
```

This table maps the lowercase and uppercase letters together into the integers 1 through 26, newline into 27, plus (+) and minus (-) into 28 and 29, and the digits into 30 through 39. Note the escape for newline. If a table is supplied, every character that is to appear either in the rules or in any valid input must be included in the table. No character may be assigned the number 0, and no character may be assigned a larger number than the size of the hardware character set.

8.19 Source Format

The general form of a lex source file is:

```
{definitions}
%%
{rules}
%%
{user subroutines}
```

The definitions section contains a combination of

1. Definitions, in the form "name space translation"
2. Included code, in the form "space code"
3. Included code, in the form

```
%{
code
%}
```

4. Start conditions, given in the form

```
%S name1 name2 ...
```

5. Character set tables, in the form

```
%T
number space character-string
%T
```

6. Changes to internal array sizes, in the form

```
%x  nnn
```

where *nnn* is a decimal integer representing an array size and *x* selects the parameter as follows:

Letter	Parameter
p	positions
n	states
e	tree nodes
a	transitions
k	packed character classes
o	output array size

Lines in the rules section have the form:

expression action

where the action may be continued on succeeding lines by using braces to delimit it.

Regular expressions in lex use the following operators:

x	The character "x"
"x"	An "x", even if x is an operator.
\x	An "x", even if x is an operator.
[xy]	The character x or y.
[x–z]	The characters x, y or z.
[^x]	Any character but x.
.	Any character but newline.
^x	An x at the beginning of a line.
<y>x	An x when lex is in start condition y.
x$	An x at the end of a line.

x?	An optional x.
x*	0,1,2, ... instances of x.
x+	1,2,3, ... instances of x.
x|y	An x or a y.
(x)	An x.
x/y	An x but only if followed by y.
{xx}	The translation of xx from the definitions section.
x{m,n}	*m* through *n* occurrences of x.

Chapter 9
Yacc: A Compiler-Compiler

9.1 Introduction 9-1

9.2 Specifications 9-4

9.3 Actions 9-6

9.4 Lexical Analysis 9-8

9.5 How the Parser Works 9-10

9.6 Ambiguity and Conflicts 9-14

9.7 Precedence 9-19

9.8 Error Handling 9-22

9.9 The Yacc Environment 9-24

9.10 Preparing Specifications 9-25

9.11 Input Style 9-25

9.12 Left Recursion 9-26

9.13 Lexical Tie-ins 9-27

9.14 Handling Reserved Words 9-27

9.15 Simulating Error and Accept in Actions 9-28

9.16 Accessing Values in Enclosing Rules 9-28

9.17 Supporting Arbitrary Value Types 9-29

9.18 A Small Desk Calculator 9-30

9.19 Yacc Input Syntax 9-32

9.20 An Advanced Example 9-34

9.21 Old Features 9-40

9.1 Introduction

Computer program input generally has some structure; every computer program that does input can be thought of as defining an input language which it accepts. An input language may be as complex as a programming language, or as simple as a sequence of numbers. Unfortunately, usual input facilities are limited, difficult to use, and often lax about checking their inputs for validity.

Yacc provides a general tool for describing the input to a computer program. The name yacc itself stands for "yet another compiler-compiler". The yacc user specifies the structures of his input, together with code to be invoked as each such structure is recognized. Yacc turns such a specification into a subroutine that handles the input process; frequently, it is convenient and appropriate to have most of the flow of control in the user's application handled by this subroutine.

The input subroutine produced by yacc calls a user-supplied routine to return the next basic input item. Thus, the user can specify his input in terms of individual input characters, or in terms of higher level constructs such as names and numbers. The user-supplied routine may also handle idiomatic features such as comment and continuation conventions, which typically defy easy grammatical specification. The class of specifications accepted is a very general one: LALR grammars with disambiguating rules.

In addition to compilers for C, APL, Pascal, RATFOR, etc., yacc has also been used for less conventional languages, including a phototypesetter language, several desk calculator languages, a document retrieval system, and a FORTRAN debugging system.

Yacc provides a general tool for imposing structure on the input to a computer program. The yacc user prepares a specification of the input process; this includes rules describing the input structure, code to be invoked when these rules are recognized, and a low-level routine to do the basic input. Yacc then generates a function to control the input process. This function, called a parser, calls the user-supplied low-level input routine (called the lexical analyzer) to pick up the basic items (called tokens) from the input stream. These tokens are organized according to the input structure rules, called grammar rules; when one of these rules has been recognized, then user code supplied for this rule, an action, is invoked; actions have the ability to return values and make use of the values of other actions.

Yacc is written in a portable dialect of C and the actions, and output subroutine, are in C as well. Moreover, many of the syntactic conventions of yacc follow C.

The heart of the input specification is a collection of grammar rules. Each rule describes an allowable structure and gives it a name. For example, one grammar rule might be:

```
date : month_name day ',' year  ;
```

Here, *date*, *month_name*, *day*, and *year* represent structures of interest in the input process; presumably, *month_name*, *day*, and *year* are defined elsewhere. The comma (,) is enclosed in single quotation marks; this implies that the comma is to appear literally in the input. The colon and semicolon merely serve as punctuation in the rule, and have no significance in controlling the input. Thus, with proper definitions, the input:

```
July 4, 1776
```

might be matched by the above rule.

An important part of the input process is carried out by the lexical analyzer. This user routine reads the input stream, recognizing the lower level structures, and communicates these tokens to the parser. A structure recognized by the lexical analyzer is called a terminal symbol, while the structure recognized by the parser is called a nonterminal symbol. To avoid confusion, terminal symbols will usually be referred to as tokens.

There is considerable leeway in deciding whether to recognize structures using the lexical analyzer or grammar rules. For example, the rules

```
month_name : 'J' 'a' 'n' ;
month_name : 'F' 'e' 'b' ;
        .
        .
        .
month_name : 'D' 'e' 'c' ;
```

might be used in the above example. The lexical analyzer would only need to recognize individual letters, and *month_name* would be a nonterminal symbol. Such low-level rules tend to waste time and space, and may complicate the specification beyond yacc's ability to deal with it. Usually, the lexical analyzer would recognize the month names, and return an indication that a *month_name* was seen; in this case, *month_name* would be a token.

Literal characters, such as the comma, must also be passed through the lexical analyzer and are considered tokens.

Specification files are very flexible. It is relatively easy to add to the above example the rule

```
date : month '/' day '/' year ;
```

allowing

```
7/4/1776
```

as a synonym for

July 4, 1776

In most cases, this new rule could be slipped in to a working system with minimal effort, and little danger of disrupting existing input.

The input being read may not conform to the specifications. These input errors are detected as early as is theoretically possible with a left-to-right scan; thus, not only is the chance of reading and computing with bad input data substantially reduced, but the bad data can usually be quickly found. Error handling, provided as part of the input specifications, permits the reentry of bad data, or the continuation of the input process after skipping over the bad data.

In some cases, yacc fails to produce a parser when given a set of specifications. For example, the specifications may be self contradictory, or they may require a more powerful recognition mechanism than that available to yacc. The former cases represent design errors; the latter cases can often be corrected by making the lexical analyzer more powerful, or by rewriting some of the grammar rules. While yacc cannot handle all possible specifications, its power compares favorably with similar systems; moreover, the constructions which are difficult for yacc to handle are also frequently difficult for human beings to handle. Some users have reported that the discipline of formulating valid yacc specifications for their input revealed errors of conception or design early in the program development.

The next several sections describe:

- The preparation of grammar rules
- The preparation of the user supplied actions associated with the grammar rules
- The preparation of lexical analyzers
- The operation of the parser
- Various reasons why yacc may be unable to produce a parser from a specification, and what to do about it.
- A simple mechanism for handling operator precedences in arithmetic expressions.
- Error detection and recovery.
- The operating environment and special features of the parsers yacc produces.
- Some suggestions which should improve the style and efficiency of the specifications.

9.2 Specifications

Names refer to either tokens or nonterminal symbols. **yacc** requires token names to be declared as such. In addition, for reasons discussed later, it is often desirable to include the lexical analyzer as part of the specification file. It may be useful to include other programs as well. Thus, every specification file consists of three sections: the declarations, (grammar) rules, and programs. The sections are separated by double percent %% marks. (The percent sign (%) is generally used in **yacc** specifications as an escape character.)

In other words, a full specification file looks like

```
declarations
%%
rules
%%
programs
```

The declaration section may be empty. Moreover, if the programs section is omitted, the second %% mark may be omitted also; thus, the smallest legal **yacc** specification is

```
%%
rules
```

Blanks, tabs, and newlines are ignored except that they may not appear in names or multicharacter reserved symbols. Comments may appear wherever a name is legal; they are enclosed in /* ... */, as in C.

The rules section is made up of one or more grammar rules. A grammar rule has the form:

```
A : BODY ;
```

A represents a nonterminal name, and *BODY* represents a sequence of zero or more names and literals. The colon and the semicolon are **yacc** punctuation.

Names may be of arbitrary length, and may be made up of letters, dot (.), the underscore (_), and noninitial digits. Uppercase and lowercase letters are distinct. The names used in the body of a grammar rule may represent tokens or nonterminal symbols.

A literal consists of a character enclosed in single quotation marks ('). As in C, the backslash (\) is an escape character within literals, and all the C escapes are recognized. Thus

'\n'	Newline
'\r'	Return
'\"'	Single quotation mark
'\\'	Backslash
'\t'	Tab
'\b'	Backspace
'\f'	Form feed
'\xxx'	"xxx" in octal

For a number of technical reasons, the ASCII NUL character ('\0' or 0) should never be used in grammar rules.

If there are several grammar rules with the same left hand side, then the vertical bar (|) can be used to avoid rewriting the left hand side. In addition, the semicolon at the end of a rule can be dropped before a vertical bar. Thus the grammar rules

```
A : B  C  D ;
A : E  F  ;
A : G  ;
```

can be given to yacc as

```
A : B  C  D
  | E  F
  | G
  ;
```

It is not necessary that all grammar rules with the same left side appear together in the grammar rules section, although it makes the input much more readable, and easier to change.

If a nonterminal symbol matches the empty string, this can be indicated in the obvious way:

```
empty : ;
```

Names representing tokens must be declared; this is most simply done by writing

```
%token name1 name2 ...
```

in the declarations section. (See Sections 3, 5, and 6 for much more discussion). Every nonterminal symbol must appear on the left side of at least one rule.

Of all the nonterminal symbols, one, called the start symbol, has particular importance. The parser is designed to recognize the start symbol; thus, this symbol represents the largest, most general structure described by the grammar rules. By default, the start symbol is taken to be the left hand side of the first grammar rule in the rules section. It is possible, and in fact desirable, to

declare the start symbol explicitly in the declarations section using the %start keyword:

```
%start symbol
```

The end of the input to the parser is signaled by a special token, called the endmarker. If the tokens up to, but not including, the endmarker form a structure which matches the start symbol, the parser function returns to its caller after the endmarker is seen; it accepts the input. If the endmarker is seen in any other context, it is an error.

It is the job of the user-supplied lexical analyzer to return the endmarker when appropriate; see section 3, below. Usually the endmarker represents some reasonably obvious I/O status, such as the end of the file or end of the record.

9.3 Actions

With each grammar rule, the user may associate actions to be performed each time the rule is recognized in the input process. These actions may return values, and may obtain the values returned by previous actions. Moreover, the lexical analyzer can return values for tokens, if desired.

An action is an arbitrary C statement, and as such can do input and output, call subprograms, and alter external vectors and variables. An action is specified by one or more statements, enclosed in curly braces { and }. For example

```
A : '(' B ')'
        {       hello( 1, "abc" );  }
```

and

```
XXX : YYY ZZZ
        { printf("a message\n");
          flag = 25;}
```

are grammar rules with actions.

To facilitate easy communication between the actions and the parser, the action statements are altered slightly. The dollar sign ($) is used as a signal to yacc in this context.

To return a value, the action normally sets the pseudo-variable **$$** to some value. For example, an action that does nothing but return the value 1 is

```
{ $$ = 1; }
```

To obtain the values returned by previous actions and the lexical analyzer, the action may use the pseudo-variables **$1**, **$2**, ..., which refer to the values returned by the components of the right side of a rule, reading from left to

right. Thus, if the rule is

```
A : B C D ;
```

for example, then $2 has the value returned by C, and $3 the value returned by D.

As a more concrete example, consider the rule

```
expr : '(' expr ')' ;
```

The value returned by this rule is usually the value of the *expr* in parentheses. This can be indicated by

```
expr : '(' expr ')' { $$ = $2 ; }
```

By default, the value of a rule is the value of the first element in it ($1). Thus, grammar rules of the form

```
A : B ;
```

frequently need not have an explicit action.

In the examples above, all the actions came at the end of their rules. Sometimes, it is desirable to get control before a rule is fully parsed. Yacc permits an action to be written in the middle of a rule as well as at the end. This rule is assumed to return a value, accessible through the usual mechanism by the actions to the right of it. In turn, it may access the values returned by the symbols to its left. Thus, in the rule

```
A : B
    { $$ = 1; }
    C
    { x = $2; y = $3; }
    ;
```

the effect is to set *x* to 1, and *y* to the value returned by C.

Actions that do not terminate a rule are actually handled by yacc by manufacturing a new nonterminal symbol name, and a new rule matching this name to the empty string. The interior action is the action triggered off by recognizing this added rule. Yacc actually treats the above example as if it had been written:

```
$ACT : /* empty */
            { $$ = 1; }
    ;

A   : B $ACT C
            { x = $2;  y = $3; }
    ;
```

In many applications, output is not done directly by the actions; rather, a data structure, such as a parse tree, is constructed in memory, and transformations are applied to it before output is generated. Parse trees are particularly easy to construct, given routines to build and maintain the tree structure desired. For example, suppose there is a C function *node*, written so that the call

```
node( L, n1, n2 )
```

creates a node with label L, and descendants n1 and n2, and returns the index of the newly created node. Then parse tree can be built by supplying actions such as:

```
expr : expr '+' expr
            { $$ = node( '+', $1, $3 ); }
```

in the specification.

The user may define other variables to be used by the actions. Declarations and definitions can appear in the declarations section, enclosed in the marks %{ and %}. These declarations and definitions have global scope, so they are known to the action statements and the lexical analyzer. For example,

```
%{ int variable = 0; %}
```

could be placed in the declarations section, making *variable* accessible to all of the actions. The yacc parser uses only names beginning in *yy*; the user should avoid such names.

In these examples, all the values are integers: a discussion of values of other types will be found in a later section.

9.4 Lexical Analysis

The user must supply a lexical analyzer to read the input stream and communicate tokens (with values, if desired) to the parser. The lexical analyzer is an integer-valued function called *yylex*. The function returns an integer, called the token number, representing the kind of token read. If there is a value associated with that token, it should be assigned to the external variable *yylval*.

The parser and the lexical analyzer must agree on these token numbers in order for communication between them to take place. The numbers may be chosen

by yacc, or chosen by the user. In either case, the # *define* mechanism of C is used to allow the lexical analyzer to return these numbers symbolically. For example, suppose that the token name *DIGIT* has been defined in the declarations section of the yacc specification file. The relevant portion of the lexical analyzer might look like:

```
yylex(){
        extern int yylval;
        int c;
        ...
        c = getchar();
        ...
        switch( c ) {
                ...
        case '0':
        case '1':
          ...
        case '9':
                yylval = c-'0';
                return( DIGIT );
                ...
                }
        ...
```

The intent is to return a token number of *DIGIT*, and a value equal to the numerical value of the digit. Provided that the lexical analyzer code is placed in the programs section of the specification file, the identifier *DIGIT* will be defined as the token number associated with the token *DIGIT*.

This mechanism leads to clear, easily modified lexical analyzers; the only pitfall is the need to avoid using any token names in the grammar that are reserved or significant in C or the parser; for example, the use of token names *if* or *while* will almost certainly cause severe difficulties when the lexical analyzer is compiled. The token name *error* is reserved for error handling, and should not be used naively.

As mentioned above, the token numbers may be chosen by yacc or by the user. In the default situation, the numbers are chosen by yacc. The default token number for a literal character is the numerical value of the character in the local character set. Other names are assigned token numbers starting at 257.

To assign a token number to a token (including literals), the first appearance of the token name or literal in the declarations section can be immediately followed by a nonnegative integer. This integer is taken to be the token number of the name or literal. Names and literals not defined by this mechanism retain their default definition. It is important that all token numbers be distinct.

For historical reasons, the endmarker must have token number 0 or negative. This token number cannot be redefined by the user. Hence, all lexical analyzers should be prepared to return 0 or negative as a token number upon reaching the

end of their input.

A very useful tool for constructing lexical analyzers is lex, discussed in a previous section. These lexical analyzers are designed to work in close harmony with yacc parsers. The specifications for these lexical analyzers use regular expressions instead of grammar rules. Lex can be easily used to produce quite complicated lexical analyzers, but there remain some languages (such as FORTRAN) which do not fit any theoretical framework, and whose lexical analyzers must be crafted by hand.

9.5 How the Parser Works

Yacc turns the specification file into a C program, which parses the input according to the specification given. The algorithm used to go from the specification to the parser is complex, and will not be discussed here (see the references for more information). The parser itself, however, is relatively simple, and understanding how it works, while not strictly necessary, will nevertheless make treatment of error recovery and ambiguities much more comprehensible.

The parser produced by yacc consists of a finite state machine with a stack. The parser is also capable of reading and remembering the next input token (called the lookahead token). The current state is always the one on the top of the stack. The states of the finite state machine are given small integer labels; initially, the machine is in state 0, the stack contains only state 0, and no lookahead token has been read.

The machine has only four actions available to it, called *shift*, *reduce*, *accept*, and *error*. A move of the parser is done as follows:

1. Based on its current state, the parser decides whether it needs a lookahead token to decide what action should be done; if it needs one, and does not have one, it calls *yylex* to obtain the next token.

2. Using the current state, and the lookahead token if needed, the parser decides on its next action, and carries it out. This may result in states being pushed onto the stack, or popped off of the stack, and in the lookahead token being processed or left alone.

The shift action is the most common action the parser takes. Whenever a shift action is taken, there is always a lookahead token. For example, in state 56 there may be an action:

```
IF      shift 34
```

which says, in state 56, if the lookahead token is IF, the current state (56) is pushed down on the stack, and state 34 becomes the current state (on the top of the stack). The lookahead token is cleared.

The reduce action keeps the stack from growing without bounds. Reduce actions are appropriate when the parser has seen the right hand side of a grammar rule, and is prepared to announce that it has seen an instance of the rule, replacing the right hand side by the left hand side. It may be necessary to consult the lookahead token to decide whether to reduce, but usually it is not; in fact, the default action (represented by a .) is often a reduce action.

Reduce actions are associated with individual grammar rules. Grammar rules are also given small integer numbers, leading to some confusion. The action

```
.       reduce 18
```

refers to grammar rule 18, while the action

```
IF      shift 34
```

refers to state 34.

Suppose the rule being reduced is

```
A : x y z ;
```

The reduce action depends on the left hand symbol (A in this case), and the number of symbols on the right hand side (three in this case). To reduce, first pop off the top three states from the stack (In general, the number of states popped equals the number of symbols on the right side of the rule). In effect, these states were the ones put on the stack while recognizing *x*, *y*, and *z*, and no longer serve any useful purpose. After popping these states, a state is uncovered which was the state the parser was in before beginning to process the rule. Using this uncovered state, and the symbol on the left side of the rule, perform what is in effect a shift of A. A new state is obtained, pushed onto the stack, and parsing continues. There are significant differences between the processing of the left hand symbol and an ordinary shift of a token, however, so this action is called a goto action. In particular, the lookahead token is cleared by a shift, and is not affected by a goto. In any case, the uncovered state contains an entry such as:

```
A       goto 20
```

causing state 20 to be pushed onto the stack, and become the current state.

In effect, the reduce action turns back the clock in the parse, popping the states off the stack to go back to the state where the right hand side of the rule was first seen. The parser then behaves as if it had seen the left side at that time. If the right hand side of the rule is empty, no states are popped off of the stack: the uncovered state is in fact the current state.

The reduce action is also important in the treatment of user-supplied actions and values. When a rule is reduced, the code supplied with the rule is executed before the stack is adjusted. In addition to the stack holding the states, another

stack, running in parallel with it, holds the values returned from the lexical analyzer and the actions. When a shift takes place, the external variable *yylval* is copied onto the value stack. After the return from the user code, the reduction is carried out. When the goto action is done, the external variable *yyval* is copied onto the value stack. The pseudo-variables **$1**, **$2**, etc., refer to the value stack.

The other two parser actions are conceptually much simpler. The accept action indicates that the entire input has been seen and that it matches the specification. This action appears only when the lookahead token is the endmarker, and indicates that the parser has successfully done its job. The error action, on the other hand, represents a place where the parser can no longer continue parsing according to the specification. The input tokens it has seen, together with the lookahead token, cannot be followed by anything that would result in a legal input. The parser reports an error, and attempts to recover the situation and resume parsing: the error recovery (as opposed to the detection of error) will be in a later section.

Consider the following example:

```
%token DING DONG DELL
%%
rhyme : sound  place
      ;
sound : DING  DONG
      ;
place : DELL
      ;
```

When **yacc** is invoked with the **–v** option, a file called *y.output* is produced, with a human-readable description of the parser. The *y.output* file corresponding to the above grammar (with some statistics stripped off the end) is:

```
state 0
        $accept : _rhyme $end

        DING shift 3
        . error

        rhyme goto 1
        sound goto 2

state 1
        $accept : rhyme_$end

        $end accept
        . error

state 2
        rhyme : sound_place

        DELL shift 5
        . error

        place goto 4


state 3
        sound : DING_DONG

        DONG shift 6
        . error

state 4
        rhyme : sound place_ (1)

        . reduce 1

state 5
        place : DELL_ (3)

        . reduce 3

state 6
        sound : DING DONG_ (2)

        . reduce 2
```

Notice that, in addition to the actions for each state, there is a description of the parsing rules being processed in each state. The underscore character (_) is used to indicate what has been seen, and what is yet to come, in each rule. Suppose the input is

DING DONG DELL

It is instructive to follow the steps of the parser while processing this input.

Initially, the current state is state 0. The parser needs to refer to the input in order to decide between the actions available in state 0, so the first token, *DING*, is read, becoming the lookahead token. The action in state 0 on *DING* is ***shift 3***, so state 3 is pushed onto the stack, and the lookahead token is cleared. State 3 becomes the current state. The next token, *DONG*, is read, becoming the lookahead token. The action in state 3 on the token *DONG* is ***shift 6***, so state 6 is pushed onto the stack, and the lookahead is cleared. The stack now contains 0, 3, and 6. In state 6, without even consulting the lookahead, the parser reduces by rule 2.

sound : DING DONG

This rule has two symbols on the right hand side, so two states, 6 and 3, are popped off of the stack, uncovering state 0. Consulting the description of state 0, looking for a goto on ***sound***,

sound goto 2

is obtained; thus state 2 is pushed onto the stack, becoming the current state.

In state 2, the next token, *DELL*, must be read. The action is ***shift 5***, so state 5 is pushed onto the stack, which now has 0, 2, and 5 on it, and the lookahead token is cleared. In state 5, the only action is to reduce by rule 3. This has one symbol on the right hand side, so one state, 5, is popped off, and state 2 is uncovered. The goto in state 2 on ***place***, the left side of rule 3, is state 4. Now, the stack contains 0, 2, and 4. In state 4, the only action is to reduce by rule 1. There are two symbols on the right, so the top two states are popped off, uncovering state 0 again. In state 0, there is a goto on ***rhyme*** causing the parser to enter state 1. In state 1, the input is read; the endmarker is obtained, indicated by ***$end*** in the *y.output* file. The action in state 1 when the endmarker is seen is to accept, successfully ending the parse.

The reader is urged to consider how the parser works when confronted with such incorrect strings as *DING DONG DONG*, *DING DONG*, *DING DONG DELL DELL*, etc. A few minutes spend with this and other simple examples will probably be repaid when problems arise in more complicated contexts.

9.6 Ambiguity and Conflicts

A set of grammar rules is ambiguous if there is some input string that can be structured in two or more different ways. For example, the grammar rule

expr : expr '-' expr

is a natural way of expressing the fact that one way of forming an arithmetic

expression is to put two other expressions together with a minus sign between them. Unfortunately, this grammar rule does not completely specify the way that all complex inputs should be structured. For example, if the input is

```
expr - expr - expr
```

the rule allows this input to be structured as either

```
( expr - expr ) - expr
```

or as

```
expr - ( expr - expr )
```

(The first is called left association, the second right association).

Yacc detects such ambiguities when it is attempting to build the parser. It is instructive to consider the problem that confronts the parser when it is given an input such as

```
expr - expr - expr
```

When the parser has read the second expr, the input that it has seen:

```
expr - expr
```

matches the right side of the grammar rule above. The parser could reduce the input by applying this rule; after applying the rule; the input is reduced to *expr* (the left side of the rule). The parser would then read the final part of the input:

```
- expr
```

and again reduce. The effect of this is to take the left associative interpretation.

Alternatively, when the parser has seen

```
expr - expr
```

it could defer the immediate application of the rule, and continue reading the input until it had seen

```
expr - expr - expr
```

It could then apply the rule to the rightmost three symbols, reducing them to *expr* and leaving

```
expr - expr
```

Now the rule can be reduced once more; the effect is to take the right associative interpretation. Thus, having read

```
expr - expr
```

the parser can do two legal things, a shift or a reduction, and has no way of deciding between them. This is called a shift/reduce conflict. It may also happen that the parser has a choice of two legal reductions; this is called a reduce/reduce conflict. Note that there are never any shift/shift conflicts.

When there are shift/reduce or reduce/reduce conflicts, yacc still produces a parser. It does this by selecting one of the valid steps wherever it has a choice. A rule describing which choice to make in a given situation is called a disambiguating rule.

Yacc invokes two disambiguating rules by default:

1. In a shift/reduce conflict, the default is to do the shift.

2. In a reduce/reduce conflict, the default is to reduce by the earlier grammar rule (in the input sequence).

Rule 1 implies that reductions are deferred whenever there is a choice, in favor of shifts. Rule 2 gives the user rather crude control over the behavior of the parser in this situation, but reduce/reduce conflicts should be avoided whenever possible.

Conflicts may arise because of mistakes in input or logic, or because the grammar rules, while consistent, require a more complex parser than yacc can construct. The use of actions within rules can also cause conflicts, if the action must be done before the parser can be sure which rule is being recognized. In these cases, the application of disambiguating rules is inappropriate, and leads to an incorrect parser. For this reason, yacc always reports the number of shift/reduce and reduce/reduce conflicts resolved by Rule 1 and Rule 2.

In general, whenever it is possible to apply disambiguating rules to produce a correct parser, it is also possible to rewrite the grammar rules so that the same inputs are read but there are no conflicts. For this reason, most previous parser generators have considered conflicts to be fatal errors. Our experience has suggested that this rewriting is somewhat unnatural, and produces slower parsers; thus, yacc will produce parsers even in the presence of conflicts.

As an example of the power of disambiguating rules, consider a fragment from a programming language involving an if-then-else construction:

```
stat : IF '(' cond ')' stat
     | IF '(' cond ')' stat ELSE stat
     ;
```

In these rules, *IF* and *ELSE* are tokens, *cond* is a nonterminal symbol describing conditional (logical) expressions, and *stat* is a nonterminal symbol describing statements. The first rule will be called the simple-if rule, and the second the

if-else rule.

These two rules form an ambiguous construction, since input of the form

```
IF ( C1 ) IF ( C2 ) S1 ELSE S2
```

can be structured according to these rules in two ways:

```
IF ( C1 ) {
        IF ( C2 ) S1
        }
ELSE S2
```

or

```
IF ( C1 ) {
        IF ( C2 ) S1
        ELSE S2
        }
```

The second interpretation is the one given in most programming languages having this construct. Each *ELSE* is associated with the last *IF* immediately preceding the *ELSE*. In this example, consider the situation where the parser has seen

```
IF ( C1 ) IF ( C2 ) S1
```

and is looking at the *ELSE*. It can immediately reduce by the simple-if rule to get

```
IF ( C1 ) stat
```

and then read the remaining input,

```
ELSE S2
```

and reduce

```
IF ( C1 ) stat ELSE S2
```

by the if-else rule. This leads to the first of the above groupings of the input.

On the other hand, the *ELSE* may be shifted, *S2* read, and then the right hand portion of

```
IF ( C1 ) IF ( C2 ) S1 ELSE S2
```

can be reduced by the if-else rule to get

IF (C1) stat

which can be reduced by the simple-if rule. This leads to the second of the above groupings of the input, which is usually desired.

Once again the parser can do two valid things – there is a shift/reduce conflict. The application of disambiguating rule 1 tells the parser to shift in this case, which leads to the desired grouping.

This shift/reduce conflict arises only when there is a particular current input symbol, *ELSE*, and particular inputs already seen, such as

IF (C1) IF (C2) S1

In general, there may be many conflicts, and each one will be associated with an input symbol and a set of previously read inputs. The previously read inputs are characterized by the state of the parser.

The conflict messages of **yacc** are best understood by examining the verbose (**–v**) option output file. For example, the output corresponding to the above conflict state might be:

```
23: shift/reduce conflict (shift 45, reduce 18) on ELSE

state 23

        stat : IF ( cond ) stat_     (18)
        stat : IF ( cond ) stat_ELSE stat

        ELSE   shift 45
        .      reduce 18
```

The first line describes the conflict, giving the state and the input symbol. The ordinary state description follows, giving the grammar rules active in the state, and the parser actions. Recall that the underline marks the portion of the grammar rules which has been seen. Thus in the example, in state 23 the parser has seen input corresponding to

IF (cond) stat

and the two grammar rules shown are active at this time. The parser can do two possible things. If the input symbol is *ELSE*, it is possible to shift into state 45. State 45 will have, as part of its description, the line

stat : IF (cond) stat ELSE_stat

since the *ELSE* will have been shifted in this state. Back in state 23, the alternative action, described by "." , is to be done if the input symbol is not mentioned explicitly in the above actions; thus, in this case, if the input symbol

is not *ELSE*, the parser reduces by grammar rule 18:

```
stat : IF '(' cond ')' stat
```

Once again, notice that the numbers following shift commands refer to other states, while the numbers following reduce commands refer to grammar rule numbers. In the *y.output* file, the rule numbers are printed after those rules which can be reduced. In most one states, there will be at most reduce action possible in the state, and this will be the default command. The user who encounters unexpected shift/reduce conflicts will probably want to look at the verbose output to decide whether the default actions are appropriate. In really tough cases, the user might need to know more about the behavior and construction of the parser than can be covered here. In this case, one of the theoretical references might be consulted; the services of a local guru might also be appropriate.

9.7 Precedence

There is one common situation where the rules given above for resolving conflicts are not sufficient; this is in the parsing of arithmetic expressions. Most of the commonly used constructions for arithmetic expressions can be naturally described by the notion of precedence levels for operators, together with information about left or right associativity. It turns out that ambiguous grammars with appropriate disambiguating rules can be used to create parsers that are faster and easier to write than parsers constructed from unambiguous grammars. The basic notion is to write grammar rules of the form

```
expr : expr OP expr
```

and

```
expr : UNARY expr
```

for all binary and unary operators desired. This creates a very ambiguous grammar, with many parsing conflicts. As disambiguating rules, the user specifies the precedence, or binding strength, of all the operators, and the associativity of the binary operators. This information is sufficient to allow yacc to resolve the parsing conflicts in accordance with these rules, and construct a parser that realizes the desired precedences and associativities.

The precedences and associativities are attached to tokens in the declarations section. This is done by a series of lines beginning with a yacc keyword: %left, %right, or %nonassoc, followed by a list of tokens. All of the tokens on the same line are assumed to have the same precedence level and associativity; the lines are listed in order of increasing precedence or binding strength. Thus,

```
%left '+' '-'
%left '*' '/'
```

describes the precedence and associativity of the four arithmetic operators. Plus and minus are left associative, and have lower precedence than star and slash, which are also left associative. The keyword %right is used to describe right associative operators, and the keyword %nonassoc is used to describe operators, like the operator .LT. in FORTRAN, that may not associate with themselves; thus,

```
A .LT. B .LT. C
```

is illegal in FORTRAN, and such an operator would be described with the keyword %nonassoc in yacc. As an example of the behavior of these declarations, the description

```
%right '='
%left '+' '-'
%left '*' '/'

%%

expr  : expr '=' expr
      | expr '+' expr
      | expr '-' expr
      | expr '*' expr
      | expr '/' expr
      | NAME
      ;
```

might be used to structure the input

```
a = b = c*d - e - f*g
```

as follows:

```
a = ( b = ( ((c*d)-e) - (f*g) ) )
```

When this mechanism is used, unary operators must, in general, be given a precedence. Sometimes a unary operator and a binary operator have the same symbolic representation, but different precedences. An example is unary and binary '-'; unary minus may be given the same strength as multiplication, or even higher, while binary minus has a lower strength than multiplication. The keyword, %prec, changes the precedence level associated with a particular grammar rule. The %prec appears immediately after the body of the grammar rule, before the action or closing semicolon, and is followed by a token name or literal. It causes the precedence of the grammar rule to become that of the following token name or literal. For example, to make unary minus have the same precedence as multiplication the rules might resemble:

```
%left '+' '-'
%left '*' '/'

%%

expr  : expr '+' expr
      | expr '-' expr
      | expr '*' expr
      | expr '/' expr
      | '-' expr %prec '*'
      | NAME
      ;
```

A token declared by %left, %right, and %nonassoc need not be, but may be, declared by %token as well.

The precedences and associativities are used by yacc to resolve parsing conflicts; they give rise to disambiguating rules. Formally, the rules work as follows:

1. The precedences and associativities are recorded for those tokens and literals that have them.

2. A precedence and associativity is associated with each grammar rule; it is the precedence and associativity of the last token or literal in the body of the rule. If the %prec construction is used, it overrides this default. Some grammar rules may have no precedence and associativity associated with them.

3. When there is a reduce/reduce conflict, or there is a shift/reduce conflict and either the input symbol or the grammar rule has no precedence and associativity, then the two disambiguating rules given at the beginning of the section are used, and the conflicts are reported.

4. If there is a shift/reduce conflict, and both the grammar rule and the input character have precedence and associativity associated with them, then the conflict is resolved in favor of the action (shift or reduce) associated with the higher precedence. If the precedences are the same, then the associativity is used; left associative implies reduce, right associative implies shift, and nonassociating implies error.

Conflicts resolved by precedence are not counted in the number of shift/reduce and reduce/reduce conflicts reported by yacc. This means that mistakes in the specification of precedences may disguise errors in the input grammar; it is a good idea to be sparing with precedences, and use them in an essentially cookbook fashion, until some experience has been gained. The *y.output* file is very useful in deciding whether the parser is actually doing what was intended.

9.8 Error Handling

Error handling is an extremely difficult area, and many of the problems are semantic ones. When an error is found, for example, it may be necessary to reclaim parse tree storage, delete or alter symbol table entries, and, typically, set switches to avoid generating any further output.

It is seldom acceptable to stop all processing when an error is found. It is more useful to continue scanning the input to find further syntax errors. This leads to the problem of getting the parser restarted after an error. A general class of algorithms to perform this involves discarding a number of tokens from the input string, and attempting to adjust the parser so that input can continue.

To allow the user some control over this process, yacc provides a simple, but reasonably general feature. The token name *error* is reserved for error handling. This name can be used in grammar rules; in effect, it suggests places where errors are expected, and recovery might take place. The parser pops its stack until it enters a state where the token *error* is legal. It then behaves as if the token *error* were the current lookahead token, and performs the action encountered. The lookahead token is then reset to the token that caused the error. If no special error rules have been specified, the processing halts when an error is detected.

In order to prevent a cascade of error messages, the parser, after detecting an error, remains in error state until three tokens have been successfully read and shifted. If an error is detected when the parser is already in error state, no message is given, and the input token is quietly deleted.

As an example, a rule of the form

```
stat : error
```

would, in effect, mean that on a syntax error the parser would attempt to skip over the statement in which the error was seen. More precisely, the parser will scan ahead, looking for three tokens that might legally follow a statement, and start processing at the first of these; if the beginnings of statements are not sufficiently distinctive, it may make a false start in the middle of a statement, and end up reporting a second error where there is in fact no error.

Actions may be used with these special error rules. These actions might attempt to reinitialize tables, reclaim symbol table space, etc.

Error rules such as the above are very general, but difficult to control. Somewhat easier are rules such as

```
stat : error ';'
```

Here, when there is an error, the parser attempts to skip over the statement, but will do so by skipping to the next ';'. All tokens after the error and before

the next ';' cannot be shifted, and are discarded. When the ';' is seen, this rule will be reduced, and any cleanup action associated with it performed.

Another form of error rule arises in interactive applications, where it may be desirable to permit a line to be reentered after an error. A possible error rule might be

```
input : error '\n' { printf( "Reenter line: "); } input
                { $$ = $4;}
```

There is one potential difficulty with this approach; the parser must correctly process three input tokens before it admits that it has correctly resynchronized after the error. If the reentered line contains an error in the first two tokens, the parser deletes the offending tokens, and gives no message; this is clearly unacceptable. For this reason, there is a mechanism that can be used to force the parser to believe that an error has been fully recovered from. The statement

```
yyerrok ;
```

in an action resets the parser to its normal mode. The last example is better written

```
input : error '\n'
                { yyerrok;
                  printf( "Reenter last line: " ); }
                input
                { $$ = $4; }
                ;
```

As mentioned above, the token seen immediately after the *error* symbol is the input token at which the error was discovered. Sometimes, this is inappropriate; for example, an error recovery action might take upon itself the job of finding the correct place to resume input. In this case, the previous lookahead token must be cleared. The statement

```
yyclearin ;
```

in an action will have this effect. For example, suppose the action after error were to call some sophisticated resynchronization routine, supplied by the user, that attempted to advance the input to the beginning of the next valid statement. After this routine was called, the next token returned by yylex would presumably be the first token in a legal statement; the old, illegal token must be discarded, and the error state reset. This could be done by a rule like

```
stat : error
        { resynch();
          yyerrok ;
          yyclearin ; }
        ;
```

These mechanisms are admittedly crude, but do allow for a simple, fairly effective recovery of the parser from many errors. Moreover, the user can get control to deal with the error actions required by other portions of the program.

9.9 The Yacc Environment

When the user inputs a specification to **yacc**, the output is a file of C programs, called *y.tab.c* on most systems. The function produced by **yacc** is called *yyparse* ; it is an integer valued function. When it is called, it in turn repeatedly calls *yylex*, the lexical analyzer supplied by the user to obtain input tokens. Eventually, either an error is detected, in which case (if no error recovery is possible) *yyparse* returns the value 1, or the lexical analyzer returns the endmarker token and the parser accepts. In this case, *yyparse* returns the value 0.

The user must provide a certain amount of environment for this parser in order to obtain a working program. For example, as with every C program, a program called *main* must be defined, that eventually calls *yyparse*. In addition, a routine called *yyerror* prints a message when a syntax error is detected.

These two routines must be supplied in one form or another by the user. To ease the initial effort of using **yacc**, a library has been provided with default versions of *main* and *yyerror*. The name of this library is system dependent; on many systems the library is accessed by a **–ly** argument to the loader. To show the triviality of these default programs, the source is given below:

```
main(){
        return( yyparse() );
        }
```

and

```
# include <stdio.h>

yyerror(s) char *s; {
        fprintf( stderr, "%s\n", s );
        }
```

The argument to *yyerror* is a string containing an error message, usually the string *syntax error*. The average application will want to do better than this. Ordinarily, the program should keep track of the input line number, and print

it along with the message when a syntax error is detected. The external integer variable *yychar* contains the lookahead token number at the time the error was detected; this may be of some interest in giving better diagnostics. Since the *main* program is probably supplied by the user (to read arguments, etc.) the yacc library is useful only in small projects, or in the earliest stages of larger ones.

The external integer variable *yydebug* is normally set to 0. If it is set to a nonzero value, the parser will output a verbose description of its actions, including a discussion of which input symbols have been read, and what the parser actions are. Depending on the operating environment, it may be possible to set this variable by using a debugging system.

9.10 Preparing Specifications

This section contains miscellaneous hints on preparing efficient, easy to change, and clear specifications. The individual subsections are more or less independent.

9.11 Input Style

It is difficult to provide rules with substantial actions and still have a readable specification file.

1. Use uppercase letters for token names, lowercase letters for nonterminal names. This rule helps you to know who to blame when things go wrong.

2. Put grammar rules and actions on separate lines. This allows either to be changed without an automatic need to change the other.

3. Put all rules with the same left hand side together. Put the left hand side in only once, and let all following rules begin with a vertical bar.

4. Put a semicolon only after the last rule with a given left hand side, and put the semicolon on a separate line. This allows new rules to be easily added.

5. Indent rule bodies by two tab stops, and action bodies by three tab stops.

The examples in the text of this section follow this style (where space permits). The user must make up his own mind about these stylistic questions; the central problem, however, is to make the rules visible through the morass of action code.

9.12 Left Recursion

The algorithm used by the yacc parser encourages so-called left recursive grammar rules: rules of the form

```
name : name rest_of_rule ;
```

These rules frequently arise when writing specifications of sequences and lists:

```
list : item
     | list ',' item
     ;
```

and

```
seq : item
    | seq item
    ;
```

In each of these cases, the first rule will be reduced for the first item only, and the second rule will be reduced for the second and all succeeding items.

With right recursive rules, such as

```
seq : item
    | item  seq
    ;
```

the parser would be a bit bigger, and the items would be seen, and reduced, from right to left. More seriously, an internal stack in the parser would be in danger of overflowing if a very long sequence were read. Thus, the user should use left recursion wherever reasonable.

It is worth considering whether a sequence with zero elements has any meaning, and if so, consider writing the sequence specification with an empty rule:

```
seq : /* empty */
    | seq item
    ;
```

Once again, the first rule would always be reduced exactly once, before the first item was read, and then the second rule would be reduced once for each item read. Permitting empty sequences often leads to increased generality. However, conflicts might arise if yacc is asked to decide which empty sequence it has seen, when it hasn't seen enough to know!

9.13 Lexical Tie-ins

Some lexical decisions depend on context. For example, the lexical analyzer might want to delete blanks normally, but not within quoted strings. Or names might be entered into a symbol table in declarations, but not in expressions.

One way of handling this situation is to create a global flag that is examined by the lexical analyzer, and set by actions. For example, suppose a program consists of 0 or more declarations, followed by 0 or more statements. Consider:

```
%{
  int dflag;
%}
  ... other declarations ...

%%

prog    : decls  stats
        ;

decls   : /* empty */
                {       dflag = 1;  }
        | decls declaration
        ;

stats   : /* empty */
                {       dflag = 0;  }
        | stats statement
        ;

    ... other rules ...
```

The flag *dflag* is now 0 when reading statements, and 1 when reading declarations, except for the first token in the first statement. This token must be seen by the parser before it can tell that the declaration section has ended and the statements have begun. In many cases, this single token exception does not affect the lexical scan.

This kind of back door approach can be over done. Nevertheless, it represents a way of doing some things that are difficult to do otherwise.

9.14 Handling Reserved Words

Some programming languages permit the user to use words like *if*, which are normally reserved, as label or variable names, provided that such use does not conflict with the legal use of these names in the programming language. This is extremely hard to do in the framework of yacc; it is difficult to pass information to the lexical analyzer telling it "this instance of 'if' is a keyword,

and that instance is a variable". The user can make a stab at it, but it is difficult. It is best that keywords be reserved; that is, be forbidden for use as variable names.

9.15 Simulating Error and Accept in Actions

The parsing actions of error and accept can be simulated in an action by use of macros *YYACCEPT* and *YYERROR*. *YYACCEPT* causes ***yyparse*** to return the value 0; *YYERROR* causes the parser to behave as if the current input symbol had been a syntax error; ***yyerror*** is called, and error recovery takes place. These mechanisms can be used to simulate parsers with multiple endmarkers or context-sensitive syntax checking.

9.16 Accessing Values in Enclosing Rules

An action may refer to values returned by actions to the left of the current rule. The mechanism is simply the same as with ordinary actions, a dollar sign followed by a digit, but in this case the digit may be 0 or negative. Consider

```
sent    : adj  noun  verb  adj  noun
              { look at the sentence ... }
        ;

adj     : THE   { $$ = THE; }
        | YOUNG { $$ = YOUNG; }
        ...
        ;

noun    : DOG   { $$ = DOG; }
        | CRONE { if( $0 == YOUNG ){
                      printf( "what?\n" );
                      }
                  $$ = CRONE;
                }
        ;
        ...
```

In the action following the word *CRONE*, a check is made preceding token shifted was not *YOUNG*. Obviously, this is only possible when a great deal is known about what might precede the symbol *noun* in the input. There is also a distinctly unstructured flavor about this. Nevertheless, at times this mechanism will save a great deal of trouble, especially when a few combinations are to be excluded from an otherwise regular structure.

9.17 Supporting Arbitrary Value Types

By default, the values returned by actions and the lexical analyzer are integers. Yacc can also support values of other types, including structures. In addition, yacc keeps track of the types, and inserts appropriate union member names so that the resulting parser will be strictly type checked. The yacc value stack is declared to be a *union* of the various types of values desired. The user declares the union, and associates union member names to each token and nonterminal symbol having a value. When the value is referenced through a **$$** or **$n** construction, yacc will automatically insert the appropriate union name, so that no unwanted conversions will take place. In addition, type checking commands such as **lint**(C) will be far more silent.

There are three mechanisms used to provide for this typing. First, there is a way of defining the union; this must be done by the user since other programs, notably the lexical analyzer, must know about the union member names. Second, there is a way of associating a union member name with tokens and nonterminals. Finally, there is a mechanism for describing the type of those few values where yacc cannot easily determine the type.

To declare the union, the user includes in the declaration section:

```
%union {
        body of union ...
        }
```

This declares the yacc value stack, and the external variables *yylval* and *yyval*, to have type equal to this union. If yacc was invoked with the **–d** option, the union declaration is copied onto the *y.tab.h* file. Alternatively, the union may be declared in a header file, and a typedef used to define the variable *YYSTYPE* to represent this union. Thus, the header file might also have said:

```
typedef union {
        body of union ...
        } YYSTYPE;
```

The header file must be included in the declarations section, by use of %{ and %}.

Once *YYSTYPE* is defined, the union member names must be associated with the various terminal and nonterminal names. The construction

```
< name >
```

is used to indicate a union member name. If this follows one of the keywords %token, %left, %right, and %nonassoc, the union member name is associated with the tokens listed. Thus, saying

```
%left <optype> '+' '-'
```

will cause any reference to values returned by these two tokens to be tagged with the union member name *optype*. Another keyword, %type, is used similarly to associate union member names with nonterminals. Thus, one might say

```
%type <nodetype> expr stat
```

There remain a couple of cases where these mechanisms are insufficient. If there is an action within a rule, the value returned by this action has no predefined type. Similarly, reference to left context values (such as $0 – see the previous subsection) leaves yacc with no easy way of knowing the type. In this case, a type can be imposed on the reference by inserting a union member name, between < and >, immediately after the first $. An example of this usage is

```
rule : aaa { $<intval>$ = 3; } bbb
            { fun( $<intval>2, $<other>0 ); }
     ;
```

This syntax has little to recommend it, but the situation arises rarely.

A sample specification is given in a later section. The facilities in this subsection are not triggered until they are used: in particular, the use of %type will turn on these mechanisms. When they are used, there is a fairly strict level of checking. For example, use of $n or $$ to refer to something with no defined type is diagnosed. If these facilities are not triggered, the yacc value stack is used to hold *int*'s, as was true historically.

9.18 A Small Desk Calculator

This example gives the complete yacc specification for a small desk calculator: the desk calculator has 26 registers, labeled *a* through *z*, and accepts arithmetic expressions made up of the operators +, –, *, /, % (mod operator), & (bitwise and), | (bitwise or), and assignment. If an expression at the top level is an assignment, the value is not printed; otherwise it is. As in C, an integer that begins with 0 (zero) is assumed to be octal; otherwise, it is assumed to be decimal.

As an example of a yacc specification, the desk calculator does a reasonable job of showing how precedences and ambiguities are used, and demonstrating simple error recovery. The major oversimplifications are that the lexical analysis phase is much simpler than for most applications, and the output is produced immediately, line by line. Note the way that decimal and octal integers are read in by the grammar rules; This job is probably better done by the lexical analyzer.

```
%{
# include <stdio.h>
# include <ctype.h>

int regs[26];
int base;

%}

%start list

%token DIGIT LETTER

%left '|'
%left '&'
%left '+' '-'
%left '*' '/' '%'
%left UMINUS /* precedence for unary minus */

%%    /* beginning of rules section */

list : /* empty */
     | list stat '\n'
     | list error '\n'
          { yyerrok; }
     ;

stat : expr
          { printf( "%d\n", $1 ); }
     |    LETTER '=' expr
               { regs[$1] = $3; }
     ;

expr : '(' expr ')'
          { $$ = $2; }
     | expr '+' expr
          { $$ = $1 + $3; }
     | expr '-' expr
          { $$ = $1 - $3; }
     | expr '*' expr
          { $$ = $1 * $3; }
     | expr '/' expr
          { $$ = $1 / $3; }
     | expr '%' expr
          { $$ = $1 % $3; }
     | expr '&' expr
          { $$ = $1 & $3; }
     | expr '|' expr
          { $$ = $1 | $3; }
```

```
        | '-' expr %prec UMINUS
            { $$ = - $2; }
        | LETTER
            { $$ = regs[$1]; }
        | number
        ;

number  : DIGIT
            { $$ = $1; base = ($1==0) ? 8 : 10; }
        | number DIGIT
            { $$ = base * $1 + $2; }
        ;

%%      /* start of programs */

yylex() {     /*  lexical analysis routine  */
              /*  returns LETTER for a lowercase letter, */
              /*  yylval = 0 through 25  */
              /*  return DIGIT for a digit, */
              /*  yylval = 0 through 9  */
              /*  all other characters */
              /*  are returned immediately  */

        int  c;

        while( (c=getchar()) == ' ' )  { /* skip blanks */ }

        /* c is now nonblank */

        if( islower( c ) ) {
              yylval = c - 'a';
              return ( LETTER );
              }
        if( isdigit( c ) ) {
              yylval = c - '0';
              return(  DIGIT  );
              }
        return( c );
        }
```

9.19 Yacc Input Syntax

This section has a description of the yacc input syntax, as a yacc specification. Context dependencies, etc., are not considered. Ironically, the yacc input specification language is most naturally specified as an LR(2) grammar; the sticky part comes when an identifier is seen in a rule, immediately following an action. If this identifier is followed by a colon, it is the start of the next rule; otherwise it is a continuation of the current rule, which just happens to have an

action embedded in it. As implemented, the lexical analyzer looks ahead after seeing an identifier, and decide whether the next token (skipping blanks, newlines, comments, etc.) is a colon. If so, it returns the token *C_IDENTIFIER*. Otherwise, it returns *IDENTIFIER*. Literals (quoted strings) are also returned as *IDENTIFIER*, but never as part of *C_IDENTIFIER*.

```
        /* grammar for the input to Yacc */

     /* basic entities */
%token IDENTIFIER      /* includes identifiers and literals */
%token C_IDENTIFIER    /* identifier followed by colon      */
%token NUMBER          /* [0-9]+    */

     /* reserved words:  %type => TYPE, %left => LEFT, etc. */

%token LEFT RIGHT NONASSOC TOKEN PREC TYPE START UNION

%token MARK    /* the %% mark */
%token LCURL   /* the %{ mark */
%token RCURL   /* the %} mark */

     /* ascii character literals stand for themselves */

%start spec

%%

spec    : defs MARK rules tail
        ;

tail    : MARK    { Eat up the rest of the file }
        | /* empty: the second MARK is optional */
        ;

defs    : /* empty */
        | defs def
        ;

def     : START IDENTIFIER
        | UNION { Copy union definition to output }
        | LCURL { Copy C code to output file } RCURL
        | ndefs rword tag nlist
        ;

rword   : TOKEN
        | LEFT
        | RIGHT
        | NONASSOC
```

```
        | TYPE
        ;

tag     : /* empty: union tag is optional */
        | '<' IDENTIFIER '>'
        ;

nlist   : nmno
        | nlist nmno
        | nlist ',' nmno
        ;

nmno    : IDENTIFIER            /* Literal illegal with %type */
        | IDENTIFIER NUMBER     /* Illegal with %type */
        ;

        /* rules section */

rules   : C_IDENTIFIER rbody prec
        | rules rule
        ;

rule    : C_IDENTIFIER rbody prec
        | '|' rbody prec
        ;

rbody   : /* empty */
        | rbody IDENTIFIER
        | rbody act
        ;

act     : '{' { Copy action, translate $$, etc. } '}'
        ;

prec    : /* empty */
        | PREC IDENTIFIER
        | PREC IDENTIFIER act
        | prec ';'
        ;
```

9.20 An Advanced Example

This section gives an example of a grammar using some of the advanced features discussed in earlier sections. The desk calculator example is modified to provide a desk calculator that does floating point interval arithmetic. The calculator understands floating point constants, the arithmetic operations +, −, *, /, unary −, and = (assignment), and has 26 floating point variables, *a* through *z*. Moreover, it also understands intervals, written

(x , y)

where x is less than or equal to y. There are 26 interval valued variables A through Z that may also be used. Assignments return no value, and print nothing, while expressions print the (floating or interval) value.

This example explores a number of interesting features of yacc and C. Intervals are represented by a structure, consisting of the left and right endpoint values, stored as a double precision values. This structure is given a type name, *INTERVAL*, by using *typedef*. The yacc value stack can also contain floating point scalars, and integers (used to index into the arrays holding the variable values). Notice that this entire strategy depends strongly on being able to assign structures and unions in C. In fact, many of the actions call functions that return structures as well.

It is also worth noting the use of *YYERROR* to handle error conditions: division by an interval containing 0, and an interval presented in the wrong order. In effect, the error recovery mechanism of yacc is used to throw away the rest of the offending line.

In addition to the mixing of types on the value stack, this grammar also demonstrates an interesting use of syntax to keep track of the type (e.g., scalar or interval) of intermediate expressions. Note that a scalar can be automatically promoted to an interval if the context demands an interval value. This causes a large number of conflicts when the grammar is run through yacc: 18 Shift/Reduce and 26 Reduce/Reduce. The problem can be seen by looking at the two input lines:

2.5 + (3.5 - 4.)

and

2.5 + (3.5 , 4.)

Notice that the 2.5 is to be used in an interval valued expression in the second example, but this fact is not known until the comma (,) is read; by this time, 2.5 is finished, and the parser cannot go back and change its mind. More generally, it might be necessary to look ahead an arbitrary number of tokens to decide whether to convert a scalar to an interval. This problem is circumvented by having two rules for each binary interval valued operator: one when the left operand is a scalar, and one when the left operand is an interval. In the second case, the right operand must be an interval, so the conversion will be applied automatically. However, there are still many cases where the conversion may be applied or not, leading to the above conflicts. They are resolved by listing the rules that yield scalars first in the specification file; in this way, the conflicts will be resolved in the direction of keeping scalar valued expressions scalar valued until they are forced to become intervals.

This way of handling multiple types is very instructive, but not very general. If there were many kinds of expression types, instead of just two, the number of

rules needed would increase dramatically, and the conflicts even more dramatically. Thus, while this example is instructive, it is better practice in a more normal programming language environment to keep the type information as part of the value, and not as part of the grammar.

Finally, a word about the lexical analysis. The only unusual feature is the treatment of floating point constants. The C library routine *atof* is used to do the actual conversion from a character string to a double precision value. If the lexical analyzer detects an error, it responds by returning a token that is illegal in the grammar, provoking a syntax error in the parser, and thence error recovery.

```
%{

# include <stdio.h>
# include <ctype.h>

typedef struct interval {
      double lo, hi;
      } INTERVAL;

INTERVAL vmul(), vdiv();

double    atof();

double    dreg[ 26 ];
INTERVAL  vreg[ 26 ];

%}

%start    lines

%union {
      int ival;
      double dval;
      INTERVAL vval;
      }

%token <ival> DREG VREG   /* indices into dreg, vreg arrays */

%token <dval> CONST       /* floating point constant */

%type  <dval> dexp        /* expression */

%type  <vval> vexp        /* interval expression */

      /* precedence information about the operators */

%left  '+' '-'
```

```
%left  '*' '/'
%left  UMINUS      /* precedence for unary minus */

%%

lines  : /* empty */
       | lines line
       ;

line   : dexp '\n'
              { printf( "%15.8f\n", $1 ); }
       | vexp '\n'
              { printf( "(%15.8f, %15.8f )\n", $1.lo, $1.hi ); }
       | DREG '=' dexp '\n'
              { dreg[$1] = $3; }
       | VREG '=' vexp '\n'
              { vreg[$1] = $3; }
       | error '\n'
              { yyerrok; }
       ;

dexp   : CONST
       | DREG
              { $$ = dreg[$1]; }
       | dexp '+' dexp
              { $$ = $1 + $3; }
       | dexp '-' dexp
              { $$ = $1 - $3; }
       | dexp '*' dexp
              { $$ = $1 * $3; }
       | dexp '/' dexp
              { $$ = $1 / $3; }
       | '-' dexp %prec UMINUS
              { $$ = - $2; }
       | '(' dexp ')'
              { $$ = $2; }
       ;

vexp   : dexp
              { $$.hi = $$.lo = $1; }
       | '(' dexp ',' dexp ')'
              {
                $$.lo  =  $2;
                $$.hi  =  $4;
                if( $$.lo  >  $$.hi ){
                      printf("interval out of order\n");
                      YYERROR;
                      }
                }
       | VREG
```

```
                { $$ = vreg[$1]; }
        | vexp '+' vexp
                { $$.hi = $1.hi + $3.hi;
                  $$.lo = $1.lo + $3.lo; }
        | dexp '+' vexp
                { $$.hi = $1 + $3.hi;
                  $$.lo = $1 + $3.lo; }
        | vexp '-' vexp
                { $$.hi = $1.hi - $3.lo;
                  $$.lo = $1.lo - $3.hi; }
        | dexp '-' vexp
                { $$.hi = $1 - $3.lo;
                  $$.lo = $1 - $3.hi;}
        | vexp '*' vexp
                { $$ = vmul( $1.lo, $1.hi, $3 ); }
        | dexp '*' vexp
                { $$ = vmul( $1, $1, $3 ); }
        | vexp '/' vexp
                { if ( dcheck( $3 ) ) YYERROR;
                  $$ = vdiv( $1.lo, $1.hi, $3 ); }
        | dexp '/' vexp
                { if ( dcheck( $3 ) ) YYERROR;
                  $$ = vdiv( $1, $1, $3 ); }
        | '-' vexp  %prec UMINUS
                { $$.hi = -$2.lo; $$.lo = -$2.hi; }
        | '(' vexp ')'
                {       $$ =  $2;  }
        ;

%%

# define BSZ 50  /* buffer size for fp numbers */

        /* lexical analysis */

yylex(){
        register  c;
                        { /* skip over blanks */ }
        while( ( c = getchar() ) == ' ' )

        if ( isupper(c) ){
                yylval.ival = c - 'A';
                return( VREG );
                }
        if ( islower(c) ){
                yylval.ival = c - 'a';
                return( DREG );
                }

        if( isdigit( c ) || c=='.' ){
```

```
        /* gobble up digits, points, exponents */

        char buf[BSZ+1], *cp = buf;
        int  dot = 0,  exp = 0;

        for( ; (cp-buf)<BSZ ; ++cp,c=getchar() ){

                *cp = c;
                if ( isdigit(c) ) continue;
                if ( c == '.' ) {
                        if ( dot++ || exp ) return( '.' );
                                        /* above causes syntax error */
                        continue;
                        }

                if (c == 'e' ) {
                        if ( exp++ ) return( 'e' );
                                        /* above causes syntax error */
                        continue;
                        }

                /* end of number */
                break;
                }
        *cp = '\0';
        if( (cp-buf) >= BSZ )
                        printf( "constant too long:  truncated\n");
        else  ungetc( c, stdin );
                        /* above pushes back last char read */
        yylval.dval = atof ( buf );
        return( CONST );
        }
return( c );
}

INTERVAL  hilo( a, b, c, d )  double  a, b, c, d; {
        /* returns the smallest interval containing a, b, c, and  d */
        /* used by *, / routines */
        INTERVAL v;

        if( a>b ) { v.hi = a; v.lo = b; }
        else { v.hi = b; v.lo = a; }

        if( c>d ) {
                if ( c>v.hi ) v.hi = c;
                if ( d<v.lo ) v.lo = d;
                }
        else  {
                if ( d>v.hi ) v.hi = d;
                if ( c<v.lo ) v.lo = c;
```

```
        }
    return( v );
    }

INTERVAL vmul( a, b, v ) double a, b; INTERVAL v; {
    return( hilo( a*v.hi, a*v.lo, b*v.hi, b*v.lo ) );
    }

dcheck( v ) INTERVAL v; {
    if( v.hi >= 0. && v.lo <= 0. ){
        printf( "divisor interval contains 0.\n" );
        return(1);
        }
    return(0);
    }

INTERVAL vdiv( a, b, v ) double a, b; INTERVAL v; {
    return( hilo( a/v.hi, a/v.lo, b/v.hi, b/v.lo ) );
    }
```

9.21 Old Features

This section mentions synonyms and features which are supported for historical continuity, but, for various reasons, are not encouraged.

1. Literals may also be delimited by double quotation marks (").

2. Literals may be more than one character long. If all the characters are alphabetic, numeric, or underscore, the type number of the literal is defined, just as if the literal did not have the quotation marks around it. Otherwise, it is difficult to find the value for such literals. The use of multicharacter literals is likely to mislead those unfamiliar with yacc, since it suggests that yacc is doing a job that must be actually done by the lexical analyzer.

3. Most places where '%' is legal, backslash (\) may be used. In particular, the double backslash (\\) is the same as %%, *left* the same as *%left*, etc.

4. There are a number of other synonyms:

   %< is the same as %left
   %> is the same as %right
   %binary and %2 are the same as %nonassoc
   %0 and %term are the same as %token
   %= is the same as %prec

5. Actions may also have the form

   ={ ... }

   and the curly braces can be dropped if the action is a single C statement.

6. C code between %{ and %} used to be permitted at the head of the rules section, as well as in the declaration section.

Appendix A
The C-Shell

A.1 Introduction A-1

A.2 Invoking the C-shell A-1

A.3 Using Shell Variables A-2

A.4 Using the C-Shell History List A-4

A.5 Using Aliases A-7

A.6 Redirecting Input and Output A-8

A.7 Creating Background and Foreground Jobs A-9

A.8 Using Built-In Commands A-9

A.9 Creating Command Scripts A-11

A.10 Using the argv Variable A-11

A.11 Substituting Shell Variables A-12

A.12 Using Expressions A-14

A.13 Using the C-Shell: A Sample Script A-15

A.14 Using Other Control Structures A-18

A.15 Supplying Input to Commands A-19

A.16 Catching Interrupts A-20

A.17 Using Other Features A-20

A.18 Starting a Loop at a Terminal A-21

A.19 Using Braces with Arguments A-22

A.20 Substituting Commands A-22

A.21 Special Characters A-23

A.1 Introduction

The C-shell program, *csh,* is a command language interpreter for XENIX system users. The C-shell, like the standard XENIX shell *sh*, is an interface between you and the XENIX commands and programs. It translates command lines typed at a terminal into corresponding system actions, gives you access to information, such as your login name, home directory, and mailbox, and lets you construct of shell procedures for automating system tasks.

This appendix explains how to use the C-shell. It also explains the syntax and function of C-shell commands and features, and shows how to use these features to create shell procedures. The C-shell is fully described in *csh*(CP) in the XENIX *Reference Manual.*

A.2 Invoking the C-shell

You can invoke the C-shell from another shell by using the **csh** command. To invoke the C-shell, type:

```
csh
```

at the standard shell's command line. You can also direct the system to invoke the C-shell for you when you log in. If you have given the C-shell as your login shell in your */etc/passwd* file entry, the system automatically starts the shell when you log in.

After the system starts the C-shell, the shell searches your home directory for the command files *.cshrc* and *.login*. If the shell finds the files, it executes the commands contained in them, then displays the C-shell prompt.

The *cshrc* file typically contains the commands you wish to execute each time you start a C-shell, and the *.login* file contains the commands you wish to execute after logging in to the system. For example, the following is the contents of a typical *.login* file:

```
set ignoreeof
set mail=(/usr/spool/mail/bill)
set time=15
set history=10
mail
```

This file contains several **set** commands. The **set** command is executed directly by the C-shell; there is no corresponding XENIX program for this command. **Set** sets the C-shell variable "ignoreeof" which shields the C-shell from logging out if CNTRL-D is hit. Instead of CNTRL-D, the **logout** command is used to log out of the system. By setting the "mail" variable, the C-shell is notified that it is to watch for incoming mail and notify you if new mail arrives.

Next the C-shell variable "time" is set to 15 causing the C-shell to automatically print out statistics lines for commands that execute for at least 15 seconds of CPU time. The variable "history" is set to 10 indicating that the C-shell will remember the last 10 commands typed in its history list, (described later).

Finally, the XENIX *mail* program is invoked.

When the C-shell finishes processing the *.login* file, it begins reading commands from the terminal, prompting for each with:

```
%
```

When you log out (by giving the **logout** command) the C-shell prints

```
logout
```

and executes commands from the file *.logout* if it exists in your home directory. After that, the C-shell terminates and XENIX logs you off the system.

A.3 Using Shell Variables

The C-shell maintains a set of variables. For example, in the above discussion, the variables "history" and "time" had the values 10 and 15. Each C-shell variable has as its value an array of zero or more strings. C-shell variables may be assigned values by the **set** command, which has several forms, the most useful of which is:

```
set name = value
```

C-shell variables may be used to store values that are to be used later in commands through a substitution mechanism. The C-shell variables most commonly referenced are, however, those that the C-shell itself refers to. By changing the values of these variables you can directly affect the behavior of the C-shell.

One of the most important variables is "path". This variable contains a list of directory names. When you type a command name at your terminal, the C-shell examines each named directory in turn, until it finds an executable file whose name corresponds to the name you typed. The **set** command with no arguments displays the values of all variables currently defined in the C-shell. The following example shows a typical default values:

```
argv    ()
home    /usr/bill
path    (. /bin /usr/bin)
prompt  %
shell   /bin/csh
status  0
```

This output indicates that the variable "path" begins with the current directory indicated by dot (.), then */bin*, and */usr/bin*. Your own local commands may be in the current directory. Normal XENIX commands reside in */bin* and */usr/bin*.

Sometimes a number of locally developed programs reside in the directory */usr/local*. If you want all C-shells that you invoke to have access to these new programs, place the command

```
set path=(. /bin /usr/bin /usr/local)
```

in the *.cshrc* file in your home directory. Try doing this, then logging out and back in. Type

```
set
```

to see that the value assigned to "path" has changed.

You should be aware that when you log in the C-shell examines each directory that you insert into your path and determines which commands are contained there, except for the current directory which the C-shell treats specially, This means that if commands are added to a directory in your search path after you have started the C-shell, they will not necessarily be found. If you wish to use a command which has been added after you have logged in, you should give the command

```
rehash
```

to the C-shell. **Rehash** causes the shell to recompute its internal table of command locations, so that it will find the newly added command. Since the C-shell has to look in the current directory on each command anyway, placing it at the end of the path specification usually works best and reduces overhead.

Other useful built in variables are "home" which shows your home directory, and "ignoreeof" which can be set in your *.login* file to tell the C-shell not to exit when it receives an end-of-file from a terminal. The variable "ignoreeof" is one of several variables whose value the C-shell does not care about; the C-shell is only concerned with whether these variables are set or unset. Thus, to set "ignoreeof" you simply type

```
set ignoreeof
```

and to unset it type

```
unset ignoreeof
```

Some other useful built-in C-shell variables are "noclobber" and "mail". The syntax

```
>filename
```

which redirects the standard output of a command just as in the regular shell, overwrites and destroys the previous contents of the named file. In this way, you may accidentally overwrite a file which is valuable. If you prefer that the C-shell not overwrite files in this way you can

```
set noclobber
```

in your *.login* file. Then typing

```
date > now
```

causes an error message if the file *now* already exists. You can type

```
date >! now
```

if you really want to overwrite the contents of *now*. The ">!" is a special syntax indicating that overwriting or "clobbering" the file is ok. (The space between the exclamation point (!) and the word "now" is critical here, as "!now" would be an invocation of the history mechanism, described below, and have a totally different effect.)

A.4 Using the C-Shell History List

The C-shell can maintain a history list into which it places the text of previous commands. It is possible to use a notation that reuses commands, or words from commands, in forming new commands. This mechanism can be used to repeat previous commands or to correct minor typing mistakes in commands.

The following figure gives a sample session involving typical usage of the history mechanism of the C-shell. Boldface indicates user input:

```
% cat bug.c
main()

{
        printf("hello);
}
% cc !$
cc bug.c
"bug.c", line 4: newline in string or char constant
"bug.c", line 5: syntax error
% ed !$
ed bug.c
29
4s/);/"&/p
      printf("hello");
w
30
q
% !c
cc bug.c
% a.out
hello% !e
ed bug.c
30
4s/lo/lo\\n/p
      printf("hello\n");
w
32
q
% !c -o bug
cc bug.c -o bug
% size a.out bug
a.out: 2784+364+1028 = 4176b = 0x1050b
bug: 2784+364+1028 = 4176b = 0x1050b
% ls -l !*
ls -l a.out bug
-rwxr-xr-x 1 bill       3932 Dec 19 09:41 a.out
-rwxr-xr-x 1 bill       3932 Dec 19 09:42 bug
% bug
hello
% pr bug.c | lpt
lpt: Command not found.
% ^lpt^lpr
pr bug.c | lpr
%
```

In this example, we have a very simple C program that has a bug or two in the file *bug.c*, which we **cat** out on our terminal. We then try to run the C compiler on it, referring to the file again as "!$", meaning the last argument to the previous command. Here the exclamation mark (!) is the history mechanism

invocation metacharacter, and the dollar sign (**$**) stands for the last argument, by analogy to the dollar sign in the editor which stands for the end-of-line. The C-shell echoed the command, as it would have been typed without use of the history mechanism, and then executed the command. The compilation yielded error diagnostics, so we now edit the file we were trying to compile, fix the bug, and run the C compiler again, this time referring to this command simply as "!c", which repeats the last command that started with the letter "c". If there were other commands beginning with the letter "c" executed recently, we could have said "!cc" or even "!cc:p" which prints the last command starting with "cc" without executing it, so that you can check to see whether you really want to execute a given command.

After this recompilation, we ran the resulting *a.out* file, and then noting that there still was a bug, ran the editor again. After fixing the program we ran the C compiler again, but tacked onto the command an extra "–o bug" telling the compiler to place the resultant binary in the file *bug* rather than *a.out*. In general, the history mechanisms may be used anywhere in the formation of new commands, and other characters may be placed before and after the substituted commands.

We then ran the **size** command to see how large the binary program images we have created were, and then we ran an "ls –l" command with the same argument list, denoting the argument list:

```
!*
```

Finally, we ran the program *bug* to see that its output is indeed correct.

To make a listing of the program, we ran the **pr** command on the file *bug.c*. In order to print the listing at a lineprinter we piped the output to **lpr**, but misspelled it as "lpt". To correct this we used a C-shell substitute, placing the old text and new text between caret (^) characters. This is similar to the substitute command in the editor. Finally, we repeated the same command with

```
!!
```

and sent its output to the lineprinter.

There are other mechanisms available for repeating commands. The **history** command prints out a numbered list of previous commands. You can then refer to these commands by number. There is a way to refer to a previous command by searching for a string which appeared in it, and there are other, less useful, ways to select arguments to include in a new command. A complete description of all these mechanisms is given in *csh*(CP) the XENIX *Reference Manual*.

A.5 Using Aliases

The C-shell has an alias mechanism that can be used to make transformations on commands immediately after they are input. This mechanism can be used to simplify the commands you type, to supply default arguments to commands, or to perform transformations on commands and their arguments. The alias facility is similar to a macro facility. Some of the features obtained by aliasing can be obtained also using C-shell command files, but these take place in another instance of the C-shell and cannot directly affect the current C-shell's environment or involve commands such as **cd** which must be done in the current C-shell.

For example, suppose there is a new version of the mail program on the system called *newmail* that you wish to use instead of the standard mail program *mail*. If you place the C-shell command

```
alias mail newmail
```

in your *.cshrc* file, the C-shell will transform an input line of the form

```
mail bill
```

into a call on *newmail*. Suppose you wish the command **ls** to always show sizes of files, that is, to always use the –s option. In this case, you can use the **alias** command to do

```
alias ls ls –s
```

or even

```
alias dir ls –s
```

creating a new command named **dir**. If we then type

```
dir ~bill
```

the C-shell translates this to

```
ls –s /usr/bill
```

Note that the tilde (~) is a special C-shell symbol that represents the user's home directory.

Thus the **alias** command can be used to provide short names for commands, to provide default arguments, and to define new short commands in terms of other commands. It is also possible to define aliases that contain multiple commands or pipelines, showing where the arguments to the original command are to be substituted using the facilities of the history mechanism. Thus the definition

```
alias cd 'cd \!* ; ls '
```

specifies an ls command after each cd command. We enclosed the entire alias definition in single quotation marks (') to prevent most substitutions from occurring and to prevent the semicolon (;) from being recognized as a metacharacter. The exclamation mark (!) is escaped with a backslash (\) to prevent it from being interpreted when the alias command is typed in. The "\!*" here substitutes the entire argument list to the prealiasing cd command; no error is given if there are no arguments. The semicolon separating commands is used here to indicate that one command is to be done and then the next. Similarly the following example defines a command that looks up its first argument in the password file.

```
alias whois 'grep \!^ /etc/passwd'
```

The C-shell currently reads the *.cshrc* file each time it starts up. If you place a large number of aliases there, C-shells will tend to start slowly. You should try to limit the number of aliases you have to a reasonable number (10 or 15 is reasonable). Too many aliases causes delays and makes the system seem sluggish when you execute commands from within an editor or other programs.

A.6 Redirecting Input and Output

In addition to the standard output, commands also have a diagnostic output that is normally directed to the terminal even when the standard output is redirected to a file or a pipe. It is occasionally useful to direct the diagnostic output along with the standard output. For instance, if you want to redirect the output of a long running command into a file and wish to have a record of any error diagnostic it produces you can type

```
command > & file
```

The "> &" here tells the C-shell to route both the diagnostic output and the standard output into *file*. Similarly you can give the command

```
command |& lpr
```

to route both standard and diagnostic output through the pipe to the lineprinter. The form

```
command >&! file
```

is used when "noclobber" is set and *file* already exists.

Finally, use the form

```
command >> file
```

to append output to the end of an existing file. If "noclobber" is set, then an

error results if *file* does not exist, otherwise the C-shell creates *file*. The form

```
command >>! file
```

lets you append to a file even if it does not exist and "noclobber" is set.

A.7 Creating Background and Foreground Jobs

When one or more commands are typed together as a pipeline or as a sequence of commands separated by semicolons, a single job is created by the C-shell consisting of these commands together as a unit. Single commands without pipes or semicolons create the simplest jobs. Usually, every line typed to the C-shell creates a job. Each of the following lines creates a job:

```
sort < data
ls -s | sort -n | head -5
mail harold
```

If the ampersand metacharacter (&) is typed at the end of the commands, then the job is started as a background job. This means that the C-shell does not wait for the job to finish, but instead, immediately prompts for another command. The job runs in the background at the same time that normal jobs, called foreground jobs, continue to be read and executed by the C-shell. Thus

```
du > usage &
```

runs the *du* program, which reports on the disk usage of your working directory, puts the output into the file *usage* and returns immediately with a prompt for the next command without waiting for *du* to finish. The *du* program continues executing in the background until it finishes, even though you can type and execute more commands in the mean time. Background jobs are unaffected by any signals from the keyboard such as the INTERRUPT or QUIT signals.

The **kill** command terminates a background job immediately. Normally, this is done by specifying the process number of the job you want killed. Process numbers can be found with the ps command.

A.8 Using Built-In Commands

This section explains how to use some of the built-in C-shell commands.

The **alias** command described above is used to assign new aliases and to display existing aliases. If given no arguments, **alias** prints the list of current aliases. It may also be given one argument, such as to show the current alias for a given string of characters. For example

```
alias ls
```

prints the current alias for the string "ls".

The **history** command displays the contents of the history list. The numbers given with the history events can be used to reference previous events that are difficult to reference contextually. There is also a C-shell variable named "prompt". By placing an exclamation point (!) in its value the C-shell will substitute the number of the current command in the history list. You can use this number to refer to a command in a history substitution. For example, you could type:

```
set prompt='\! % '
```

Note that the exclamation mark (!) had to be escaped here even within backslashes.

The **logout** command is used to terminate a login C-shell that has "ignoreeof" set.

The **rehash** command causes the C-shell to recompute a table of command locations. This is necessary if you add a command to a directory in the current C-shell's search path and want the C-shell to find it, since otherwise the hashing algorithm may tell the C-shell that the command wasn't in that directory when the hash table was computed.

The **repeat** command is used to repeat a command several times. Thus to make 5 copies of the file *one* in the file *five* you could type:

```
repeat 5 cat one >> five
```

The **setenv** command can be used to set variables in the environment. Thus

```
setenv TERM adm3a
```

sets the value of the environment variable "TERM" to "adm3a". The program *env* exists to print out the environment. For example, its output might look like this:

```
HOME=/usr/bill
SHELL=/bin/csh
PATH=:/usr/ucb:/bin:/usr/bin:/usr/local
TERM=adm3a
USER=bill
```

The **source** command is used to force the current C-shell to read commands from a file. Thus

```
source .cshrc
```

can be used after editing in a change to the *.cshrc* file that you wish to take effect before the next time you login.

The **time** command is used to cause a command to be timed no matter how much CPU time it takes. Thus

```
time cp /etc/rc /usr/bill/rc
```

displays:

```
0.0u 0.1s 0:01 8%
```

Similarly

```
time wc /etc/rc /usr/bill/rc
```

displays:

```
  52   178   1347 /etc/rc
  52   178   1347 /usr/bill/rc
 104   356   2694 total
0.1u 0.1s 0:00 13%
```

This indicates that the **cp** command used a negligible amount of user time (u) and about 1/10th of a second system time (s); the elapsed time was 1 second (0:01). The word count command **wc** used 0.1 seconds of user time and 0.1 seconds of system time in less than a second of elapsed time. The percentage "13%" indicates that over the period when it was active the **wc** command used an average of 13 percent of the available CPU cycles of the machine.

The **unalias** and **unset** commands are used to remove aliases and variable definitions from the C-shell. The command **unsetenv** removes variables from the environment.

A.9 Creating Command Scripts

It is possible to place commands in files and to cause C-shells to be invoked to read and execute commands from these files, which are called C-shell scripts. This section describes the C-shell features that are useful when creating C-shell scripts.

A.10 Using the argv Variable

A **csh** command script may be interpreted by saying

```
csh script argument ...
```

where *script* is the name of the file containing a group of C-shell commands and

argument is a sequence of command arguments. The C-shell places these arguments in the variable "argv" and then begins to read commands from *script*. These parameters are then available through the same mechanisms that are used to reference any other C-shell variables.

If you make the file *script* executable by doing

```
chmod 755 script
```

or

```
chmod +x script
```

and then place a C-shell comment at the beginning of the C-shell script (i.e., begin the file with a number sign (#)) then */bin/csh* will automatically be invoked to execute *script* when you type

```
script
```

If the file does not begin with a number sign (#) then the standard shell */bin/sh* will be used to execute it.

A.11 Substituting Shell Variables

After each input line is broken into words and history substitutions are done on it, the input line is parsed into distinct commands. Before each command is executed a mechanism know as variable substitution is performed on these words. Keyed by the dollar sign ($), this substitution replaces the names of variables by their values. Thus

```
echo $argv
```

when placed in a command script would cause the current value of the variable "argv" to be echoed to the output of the C-shell script. It is an error for "argv" to be unset at this point.

A number of notations are provided for accessing components and attributes of variables. The notation

```
$?name
```

expands to 1 if *name* is set or to 0 if *name* is not set. It is the fundamental mechanism used for checking whether particular variables have been assigned values. All other forms of reference to undefined variables cause errors.

The notation

```
$#name
```

expands to the number of elements in the variable "name". To illustrate, examine the following terminal session (input is in boldface):

```
% set argv=(a b c)
1
% echo $#argv
3
% unset argv
% echo $?argv
0
% echo $argv
Undefined variable: argv.
%
```

It is also possible to access the components of a variable that has several values. Thus

```
$argv[1]
```

gives the first component of "argv" or in the example above "a". Similarly

```
$argv[$#argv]
```

would give "c", and

```
$argv[1-2]
```

would give:

```
a b
```

Other notations useful in C-shell scripts are

```
$n
```

where *n* is an integer. This is shorthand for

```
$argv[ n ]
```

the *n*'th parameter and

```
$*
```

which is a shorthand for

```
$argv
```

The form

$$

expands to the process number of the current C-shell. Since this process number is unique in the system, it is often used in the generation of unique temporary filenames. The form

$<

is quite special and is replaced by the next line of input read from the C-shell's standard input (not the script it is reading). This is useful for writing C-shell scripts that are interactive, reading commands from the terminal, or even writing a C-shell script that acts as a filter, reading lines from its input file. Thus, the sequence

```
echo -n 'yes or no? '
set a=($<)
```

writes out the prompt

```
yes or no?
```

without a newline and then reads the answer into the variable "a". In this case "$#a" is 0 if either a blank line or CNTRL-D is typed.

One minor difference between "$*n*" and "$argv[*n*]" should be noted here. The form "$argv[*n*]" will yield an error if *n* is not in the range 1–$#argv while "$n" will never yield an out-of-range subscript error. This is for compatibility with the way older shells handle parameters.

Another important point is that it is never an error to give a subrange of the form "n–"; if there are less than "n" components of the given variable then no words are substituted. A range of the form "m–n" likewise returns an empty vector without giving an error when "m" exceeds the number of elements of the given variable, provided the subscript "n" is in range.

A.12 Using Expressions

To construct useful C-shell scripts, the C-shell must be able to evaluate expressions based on the values of variables. In fact, all the arithmetic operations of the C language are available in the C-shell with the same precedence that they have in C. In particular, the operations "==" and "!=" compare strings and the operators "&&" and "| |" implement the logical AND and OR operations. The special operators "=~" and "!~" are similar to "==" and "!=" except that the string on the right side can have pattern matching characters (like *, ? or [and]). These operators test whether the string on the left matches the pattern on the right.

The C-shell also allows file enquiries of the form

–? filename

where question mark (?) is replaced by a number of single characters. For example, the expression primitive

–e filename

tells whether *filename* exists. Other primitives test for read, write and execute access to the file, whether it is a directory, or if it has nonzero length.

It is possible to test whether a command terminates normally, by using a primitive of the form

{ *command* }

which returns 1 if the command exits normally with exit status 0, or 0 if the command terminates abnormally or with exit status nonzero. If more detailed information about the execution status of a command is required, it can be executed and the "status" variable examined in the next command. Since "$status" is set by every command, its value is always changing.

For the full list of expression components, see **csh**(CP). in the XENIX *Reference Manual.*

A.13 Using the C-Shell: A Sample Script

A sample C-shell script follows that uses the expression mechanism of the C-shell and some of its control structures:

```
#
# Copyc copies those C programs in the specified list
# to the directory ~/backup if they differ from the files
# already in ~/backup
#
set noglob
foreach i ($argv)

        if ($i !~ *.c) continue  # not a .c file so do nothing

        if (! -r ~/backup/$i:t) then
                echo $i:t not in backup... not cp\'ed
                continue
        endif

        cmp -s $i ~/backup/$i:t # to set $status

        if ($status != 0) then
                echo new backup of $i
                cp $i ~/backup/$i:t
        endif
end
```

This script uses the **foreach** command. The command executes the other commands between the **foreach** and the matching **end**. for each of the values given between parentheses with the named variable "i" which is set to successive values in the list. Within this loop we may use the command **break** to stop executing the loop and **continue** to prematurely terminate one iteration and begin the next. After the **foreach** loop the iteration variable (*i* in this case) has the value at the last iteration.

The "noglob" variable is set to prevent filename expansion of the members of "argv". This is a good idea, in general, if the arguments to a C-shell script are filenames which have already been expanded or if the arguments may contain filename expansion metacharacters. It is also possible to quote each use of a "$" variable expansion, but this is harder and less reliable.

The other control construct is a statement of the form

```
if ( expression ) then
        command
        ...
endif
```

The placement of the keywords in this statement is not flexible due to the current implementation of the C-shell. The following two formats are not acceptable to the C-shell:

```
if (expression) # Won't work!
then
        command
        ...
endif
```

and

```
if (expression) then command endif # Won't work
```

The C-shell does have another form of the if statement:

```
if ( expression ) command
```

which can be written

```
if ( expression ) \
        command
```

Here we have escaped the newline for the sake of appearance. The command must not involve "|", "&" or ";" and must not be another control command. The second form requires the final backslash (\) to immediately precede the end-of-line.

The more general **if** statements above also admit a sequence of **else–if** pairs followed by a single **else** and an **endif**, for example:

```
if ( expression ) then
        commands
else if (expression ) then
        commands
...

else
        commands
endif
```

Another important mechanism used in C-shell scripts is the colon (:) modifier. We can use the modifier :r here to extract the root of a filename or :e to extract the extension. Thus if the variable "i" has the value */mnt/foo.bar* then

```
echo $i $i:r $i:e
```

produces

```
/mnt/foo.bar /mnt/foo bar
```

This example shows how the :r modifier strips off the trailing ".bar" and the :e modifier leaves only the "bar". Other modifiers take off the last component of a pathname leaving the head :h or all but the last component of a pathname

leaving the tail :t. These modifiers are fully described in the **csh**(CP) entry in the XENIX *Reference Manual*. It is also possible to use the command substitution mechanism to perform modifications on strings to then reenter the C-shell environment. Since each usage of this mechanism involves the creation of a new process, it is much more expensive to use than the colon (:) modification mechanism. It is also important to note that the current implementation of the C-shell limits the number of colon modifiers on a "$" substitution to 1. Thus

```
% echo $i $i:h:t
```

produces

```
/a/b/c /a/b:t
```

and does not do what you might expect.

Finally, we note that the number sign character (#) lexically introduces a C-shell comment in C-shell scripts (but not from the terminal). All subsequent characters on the input line after a number sign are discarded by the C-shell. This character can be quoted using "'" or "\" to place it in an argument word.

A.14 Using Other Control Structures

The C-shell also has control structures **while** and **switch** similar to those of C. These take the forms

```
while ( expression )
        commands
end
```

and

```
switch ( word )

case str1:
        commands
        breaksw

 ...

case strn:
        commands
        breaksw

default:
        commands
        breaksw

endsw
```

For details see the manual section for **csh**(CP). C programmers should note that we use **breaksw** to exit from a **switch** while **break** exits a **while** or **foreach** loop. A common mistake to make in C-shell scripts is to use **break** rather than **breaksw** in switches.

Finally, the C-shell allows a **goto** statement, with labels looking like they do in C:

```
loop:
        commands
        goto loop
```

A.15 Supplying Input to Commands

Commands run from C-shell scripts receive by default the standard input of the C-shell which is running the script. It allows C-shell scripts to fully participate in pipelines, but mandates extra notation for commands that are to take inline data.

Thus we need a metanotation for supplying inline data to commands in C-shell scripts. For example, consider this script which runs the editor to delete leading blanks from the lines in each argument file:

```
# deblank — remove leading blanks
foreach i ($argv)
ed - $i << 'EOF'
1,$s/^[ ]*//
w
q
'EOF'
end
```

The notation

```
<< 'EOF'
```

means that the standard input for the ed command is to come from the text in the C-shell script file up to the next line consisting of exactly EOF. The fact that the EOF is enclosed in single quotation marks ('), i.e., it is quoted, causes the C-shell to not perform variable substitution on the intervening lines. In general, if any part of the word following the "<<" which the C-shell uses to terminate the text to be given to the command is quoted then these substitutions will not be performed. In this case since we used the form "1,$" in our editor script we needed to insure that this dollar sign was not variable substituted. We could also have insured this by preceding the dollar sign ($) with a backslash (\ ), i.e.:

```
1,\$s/^[ ]*//
```

Quoting the EOF terminator is a more reliable way of achieving the same thing.

A.16 Catching Interrupts

If our C-shell script creates temporary files, we may wish to catch interruptions of the C-shell script so that we can clean up these files. We can then do

```
onintr label
```

where *label* is a label in our program. If an interrupt is received the C-shell will do a "goto label" and we can remove the temporary files, then do an **exit** command (which is built in to the C-shell) to exit from the C-shell script. If we wish to exit with nonzero status we can write

```
exit (1)
```

to exit with status 1.

A.17 Using Other Features

There are other features of the C-shell useful to writers of C-shell procedures. The **verbose** and **echo** options and the related –v and –x command line

options can be used to help trace the actions of the C-shell. The –n option causes the C-shell only to read commands and not to execute them and may sometimes be of use.

One other thing to note is that the C-shell will not execute C-shell scripts that do not begin with the number sign character (#), that is C-shell scripts that do not begin with a comment.

There is also another quotation mechanism using the double quotation mark ("), which allows only some of the expansion mechanisms we have so far discussed to occur on the quoted string and serves to make this string into a single word as the single quote (´) does.

A.18 Starting a Loop at a Terminal

It is occasionally useful to use the **foreach** control structure at the terminal to aid in performing a number of similar commands. For instance, if there were three shells in use on a particular system, */bin/sh*, */bin/nsh*, and */bin/csh*, you could count the number of persons using each shell by using the following commands:

```
grep –c csh$ /etc/passwd
grep –c nsh$ /etc/passwd
grep –c –v sh$ /etc/passwd
```

Since these commands are very similar we can use **foreach** to simplify them:

```
$ foreach i (´sh$´ ´csh$´ ´–v sh$´)
? grep –c $i /etc/passwd
? end
```

Note here that the C-shell prompts for input with "? " when reading the body of the loop. This occurs only when the **foreach** command is entered interactively.

Also useful with loops are variables that contain lists of filenames or other words. For example, examine the following terminal session:

```
% set a=(`ls`)
% echo $a
csh.n csh.rm
% ls
csh.n
csh.rm
% echo $#a
2
```

The set command here gave the variable "a" a list of all the filenames in the current directory as value. We can then iterate over these names to perform

any chosen function.

The output of a command within back quotation marks (`) is converted by the C-shell to a list of words. You can also place the quoted string within double quotation marks (") to take each (nonempty) line as a component of the variable. This prevents the lines from being split into words at blanks and tabs. A modifier :x exists which can be used later to expand each component of the variable into another variable by splitting the original variable into separate words at embedded blanks and tabs.

A.19 Using Braces with Arguments

Another form of filename expansion involves the characters, "{" and "}". These characters specify that the contained strings, separated by commas (,) are to be consecutively substituted into the containing characters and the results expanded left to right. Thus

```
A{str1,str2,...strn}B
```

expands to

```
Astr1B Astr2B ... AstrnB
```

This expansion occurs before the other filename expansions, and may be applied recursively (i.e., nested). The results of each expanded string are sorted separately, left to right order being preserved. The resulting filenames are not required to exist if no other expansion mechanisms are used. This means that this mechanism can be used to generate arguments which are not filenames, but which have common parts.

A typical use of this would be

```
mkdir ~/{hdrs,retrofit,csh}
```

to make subdirectories *hdrs*, *retrofit* and *csh* in your home directory. This mechanism is most useful when the common prefix is longer than in this example:

```
chown root /usr/demo/{file1,file2,...}
```

A.20 Substituting Commands

A command enclosed in accent symbols (`) is replaced, just before filenames are expanded, by the output from that command. Thus, it is possible to do

```
set pwd=`pwd`
```

to save the current directory in the variable "pwd" or to do

```
vi `grep -l TRACE *.c`
```

to run the editor **vi** supplying as arguments those files whose names end in *.c* which have the string "TRACE" in them. Command expansion also occurs in input redirected with "<<" and within quotation marks ("). Refer to **csh**(CP) in the XENIX *Reference Manual* for more information.

A.21 Special Characters

The following table lists the special characters of **csh** and the XENIX system. A number of these characters also have special meaning in expressions. See the **csh** manual section for a complete list.

Syntactic metacharacters

;	Separates commands to be executed sequentially
|	Separates commands in a pipeline
( )	Brackets expressions and variable values
&	Follows commands to be executed without waiting for completion

Filename metacharacters

/	Separates components of a file's pathname
.	Separates root parts of a filename from extensions
?	Expansion character matching any single character
*	Expansion character matching any sequence of characters
[ ]	Expansion sequence matching any single character from a set of characters
~	Used at the beginning of a filename to indicate home directories
{ }	Used to specify groups of arguments with common parts

Quotation metacharacters

\	Prevents meta-meaning of following single character
'	Prevents meta-meaning of a group of characters

Like ´, but allows variable and command expansion

Input/output metacharacters

<	Indicates redirected input
>	Indicates redirected output

Expansion/Substitution Metacharacters

$	Indicates variable substitution
!	Indicates history substitution
:	Precedes substitution modifiers
^	Used in special forms of history substitution
`	Indicates command substitution

Other Metacharacters

#	Begins scratch filenames; indicates C-shell comments
–	Prefixes option (flag) arguments to commands

Appendix B
C Language Portability

B.1 Introduction B-1

B.2 Program Portability B-2

B.3 Machine Hardware B-2
- B.3.1 Byte Length B-2
- B.3.2 Word Length B-2
- B.3.3 Storage Alignment B-3
- B.3.4 Byte Order in a Word B-4
- B.3.5 Bitfields B-5
- B.3.6 Pointers B-5
- B.3.7 Address Space B-6
- B.3.8 Character Set B-6

B.4 Compiler Differences B-7
- B.4.1 Signed/Unsigned char, Sign Extension B-7
- B.4.2 Shift Operations B-7
- B.4.3 Identifier Length B-8
- B.4.4 Register Variables B-8
- B.4.5 Type Conversion B-8
- B.4.6 Functions With Variable Number of Arguments - B-9
- B.4.7 Side Effects, Evaluation Order B-11

B.5 Program Environment Differences B-12

B.6 Portability of Data B-12

B.7 Lint B-13

B.8 Byte Ordering Summary B-13

B.1 Introduction

The standard definition of the C programming language leaves many details to be decided by individual implementations of the language. These unspecified features of the language detract from its portability and must be studied when attempting to write portable C code.

Most of the issues affecting C portability arise from differences in either target machine hardware or compilers. C was designed to compile to efficient code for the target machine (initially a PDP-11) and so many of the language features not precisely defined are those that reflect a particular machine's hardware characteristics.

This appendix highlights the various aspects of C that may not be portable across different machines and compilers. It also briefly discusses the portability of a C program in terms of its environment, which is determined by the system calls and library routines it uses during execution, file pathnames it requires, and other items not guaranteed to be constant across different systems.

The C language has been implemented on many different computers with widely different hardware characteristics, from small 8-bit microprocessors to large mainframes. This appendix is concerned with the portability of C code in the XENIX programming environment. This is a more restricted problem to consider since all XENIX systems to date run on hardware with the following basic characteristics:

- ASCII character set
- 8-bit bytes
- 2-byte or 4-byte integers
- Two's complement arithmetic

These features are not formally defined for the language and may not be found in of all implementations of C. However, the remainder of this appendix is devoted to those systems where these basic assumptions hold.

The C language definition contains no specification of how input and output is performed. This is left to system calls and library routines on individual systems. Within XENIX systems there are system calls and library routines that can be considered portable. These are described briefly in a later section.

This appendix is not intended as a C language primer. It is assumed that the reader is familiar with C, and with the basic architecture of common microprocessors.

B.2 Program Portability

A program is portable if it can be compiled and run successfully on different machines without alteration. There are many ways to write portable programs. The first is to avoid using inherently nonportable language features. The second is to isolate any nonportable interactions with the environment, such as I/O to nonstandard devices. For example programs should avoid hard-coding pathnames unless a pathname is common to all systems (e.g., */etc/passwd*).

Files required at compiletime (i.e., include files) may also introduce nonportability if the pathnames are not the same on all machines. In some cases include files containing machine parameters can be used to make the source code itself portable.

B.3 Machine Hardware

Differences in the hardware of the various target machines and differences in the corresponding C compilers cause the greatest number of portability problems. This section lists problems commonly encountered on XENIX systems.

B.3.1 Byte Length

By definition, the **char** data type in C must be large enough hold as positive integers all members of a machine's character set. For the machines described in this appendix, the **char** size is exactly an 8 bit byte.

B.3.2 Word Length

In C, the size of the basic data types for a given implementation are not formally defined. Thus they often follow the most natural size for the underlying machine. It is safe to assume that **short** is no longer than **long**. Beyond that no assumptions are portable. For example on some machines **short** is the same length as **int**, whereas on others **long** is the same length as **int**.

Programs that need to know the size of a particular data type should avoid hard-coded constants where possible. Such information can usually be written in a fairly portable way. For example the maximum positive integer (on a two's complement machine) can be obtained with:

```
#define MAXPOS          ((int)(((unsigned)-1) >> 1))
```

This is preferable to something like:

```
#ifdef PDP11
#define MAXPOS 32767
#else
        ...
#endif
```

To find the number of bytes in an **int** use "sizeof (int)" rather than 2, 4, or some other nonportable constant.

B.3.3 Storage Alignment

The C language defines no particular layout for storage of data items relative to each other, or for storage of elements of structures or unions within the structure or union.

Some CPU's, such as the PDP-11 and M68000 require that data types longer than one byte be aligned on even byte address boundaries. Others, such as the 8086 and VAX-11 have no such hardware restriction. However, even with these machines, most compilers generate code that aligns words, structures, arrays, and long words on even addresses, or even long word addresses. Thus, on the VAX-11, the following code sequence gives "8", even though the VAX hardware can access an **int** (a 4-byte word) on any physical starting address:

```
struct s_tag {
        char c;
        int  i;
};
printf("%d\n",sizeof(struct s_tag));
```

The principal implications of this variation in data storage are that data accessed as nonprimitive data types is not portable, and code that makes use of knowledge of the layout on a particular machine is not portable.

Thus unions containing structures are nonportable if the union is used to access the same data in different ways. Unions are only likely to be portable if they are used simply to have different data in the same space at different times. For example, if the following union were used to obtain 4 bytes from a long word, the code would not be portable:

```
union {
        char c[4];
        long lw;
} u;
```

The *sizeof* operator should always be used when reading and writing structures:

```
struct s_tag st;

  ...

write(fd, &st, sizeof(st));
```

This ensures portability of the source code. It does not produce a portable data file. Portability of data is discussed in a later section.

Note that the *sizeof* operator returns the number of bytes an object would occupy in an array. Thus on machines where structures are always aligned to begin on a word boundary in memory, the *sizeof* operator will include any necessary padding for this in the return value, even if the padding occurs after all useful data in the structure. This occurs whether or not the argument is actually an array element.

B.3.4 Byte Order in a Word

The variation in byte order in a word affects the portability of data more than the portability of source code. However any program that makes use of knowledge of the internal byte order in a word is not portable. For example, on some systems there is an include file *misc.h* that contains the following structure declaration:

```
/*
* structure to access an
* integer in bytes
*/
struct {
        char    lobyte;
        char    hibyte;
};
```

With certain less restrictive compilers this could be used to access the high and low order bytes of an integer separately, and in a completely nonportable way. The correct way to do this is to use mask and shift operations to extract the required byte:

```
#define LOBYTE(i) (i & 0xff)
#define HIBYTE(i) ((i >> 8) & 0xff)
```

Note that even this operation is only applicable to machines with two bytes in an **int**.

One result of the byte ordering problem is that the following code sequence will not always perform as intended:

```
int c = 0;

read(fd, &c, 1);
```

On machines where the low order byte is stored first, the value of "c" will be the byte value read. On other machines the byte is read into some byte other than the low order one, and the value of "c" is different.

B.3.5 Bitfields

Bitfields are not implemented in all C compilers. When they are, no field may be larger than an **int**, and no field can overlap an **int** boundary. If necessary the compiler will leave gaps and move to the next **int** boundary.

The C language makes no guarantees about whether fields are assigned left to right, or right to left in an **int**. Thus, while bitfields may be useful for storing flags and other small data items, their use in unions to dissect bits from other data is definitely nonportable.

To ensure portability no individual field should exceed 16 bits.

B.3.6 Pointers

The C language is fairly generous in allowing manipulation of pointers, to the extent that most compilers will not object to nonportable pointer operations. The *lint* program is particularly useful for detecting questionable pointer assignments and comparisons.

The common nonportable use of pointers is the use of casts to assign one pointer to another pointer of a different data type. This almost always makes some assumption about the internal byte ordering and layout of the data type, and is therefore nonportable. In the following code, the byte order in the given array is not portable:

```
char c[4];
long *lp;

lp = (long *)&c[0];
*lp = 0x12345678L;
```

The *lint* program will issue warning messages about such uses of pointers. Code like this is very rarely necessary or valid. It is acceptable, however, when using the *malloc* function to allocate space for variables that do not have **char** type. The routine is declared as type **char** * and the return value is cast to the type to be stored in the allocated memory. If this type is not **char** * then *lint* will issue a warning concerning illegal type conversion. In addition, the *malloc* function is written to always return a starting address suitable for storing all types of data. *Lint* does not know this, so it gives a warning about possible data

alignment problems too. In the following example, *malloc* is used to obtain memory for an array of 50 integers.

```
extern char *malloc();
int *ip;

ip = (int *)malloc(50);
```

This example will attract a warning message from *lint*.

The C Reference manual states that a pointer can be assigned (or cast) to an integer large enough to hold it. Note that the size of the **int** type depends on the given machine and implementation. This type is a **long** on some machines and **short** on others. In general, do not assume that "sizeof(char *) == sizeof(int)".

In most implementations, the null pointer value, "NULL" is defined to be the integer value 0. This can lead to problems for functions that expect pointer arguments larger than integers. For portable code, always use

```
func( (char *)NULL );
```

to pass a "NULL" value of the correct size.

B.3.7 Address Space

The address space available to a program running under XENIX varies considerably from system to system. On a small PDP-11 there may be only 64K bytes available for program and data combined. Larger PDP-11's, and some 16 bit microprocessors allow 64K bytes of data, and 64K bytes of program text. Other machines may allow considerably more text, and possibly more data as well.

Large programs, or programs that require large data areas may have portability problems on small machines.

B.3.8 Character Set

The C language does not require the use of the ASCII character set. In fact, the only character set requirements are all characters must fit in the **char** data type, and all characters must have positive values.

In the ASCII character set, all characters have values between zero and 127. Thus they can all be represented in 7 bits, and on an 8-bits-per-byte machine are all positive, whether **char** is treated as signed or unsigned.

There is a set of macros defined under XENIX in the header file */usr/include/ctype.h* that should be used for most tests on character

quantities. They provide insulation from the internal structure of the character set and in most cases their names are more meaningful than the equivalent line of code. Compare

```
if(isupper(c))
```

to

```
if((c >= 'A') && (c <= 'Z'))
```

With some of the other macros, such as *isdigit* to test for a hex digit, the advantage is even greater. Also, the internal implementation of the macros makes them more efficient than an explicit test with an 'if' statement.

B.4 Compiler Differences

There are a number of C compilers running under XENIX. On PDP-11 systems there is the so-called "Ritchie" compiler. Also on the 11, and on most other systems, there is the Portable C Compiler.

B.4.1 Signed/Unsigned char, Sign Extension

The current state of the signed versus unsigned **char** problem is best described as unsatisfactory.

The sign extension problem is a serious barrier to writing portable C, and the best solution at present is to write defensive code that does not rely on particular implementation features.

B.4.2 Shift Operations

The left shift operator, "<<" shifts its operand a number of bits left, filling vacated bits with zero. This is a so-called logical shift. The right shift operator, ">>" when applied to an unsigned quantity, performs a logical shift operation. When applied to a signed quantity, the vacated bits may be filled with zero (logical shift) or with sign bits (arithmetic shift). The decision is implementation dependent, and code that uses knowledge of a particular implementation is nonportable.

The PDP-11 compilers use arithmetic right shift. To avoid sign extension it is necessary to shift and mask out the appropriate number of high order bits:

```
char c;

c = (c >> 3) & 0x1f;
```

You can also avoid sign extension by using using the divide operator:

```
char c;

c = c / 8;
```

B.4.3 Identifier Length

The use of long symbols and identifier names will cause portability problems with some compilers. To avoid these problems, a program should keep the following symbols as short as possible:

— C Preprocessor Symbols

— C Local Symbols

— C External Symbols

The loader used may also place a restriction on the number of unique characters in C external symbols.

Symbols unique in the first six characters are unique to most C language processors.

On some non-XENIX C implementations, uppercase and lowercase letters are not distinct in identifiers.

B.4.4 Register Variables

The number and type of register variables in a function depends on the machine hardware and the compiler. Excess and invalid register declarations are treated as nonregister declarations and should not cause a portability problem. On a PDP-11, up to three register declarations are significant, and they must be of type **int**, **char**, or pointer. While other machines and compilers may support declarations such as

```
register unsigned short
```

this should not be relied upon.

Since the compiler ignores excess variables of register type, the most important register type variables should be declared first. Thus, if any are ignored, they will be the least important ones.

B.4.5 Type Conversion

The C language has some rules for implicit type conversion; it also allows explicit type conversions by type casting. The most common portability

problem in implicit type conversion is unexpected sign extension. This is a potential problem whenever something of type **char** is compared with an **int**.

For example

```
char c;

if(c == 0x80)
        ...
```

will never evaluate true on a machine which sign extends since "c" is sign extended before the comparison with 0x80, an **int**.

The only safe comparison between **char** type and an **int** is the following:

```
char c;

if(c == 'x')
        ...
```

This is reliable because C guarantees all characters to be positive. The use of hard-coded octal constants is subject to sign extension. For example the following program prints "ff80" on a PDP-11:

```
main()
{
        printf("%x\n",'\200');
}
```

Type conversion also takes place when arguments are passed to functions. Types **char** and **short** become **int**. Machines that sign extend **char** can give surprises. For example the following program gives –128 on some machines:

```
char c = 128;
printf("%d\n",c);
```

This is because "c" is converted to **int** before passing o the function. The function itself has no knowledge of the original type of the argument, and is expecting an **int**. The correct way to handle this is to code defensively and allow for the possibility of sign extension:

```
char c = 128;
printf("%d\n", c & 0xff);
```

B.4.6 Functions With Variable Number of Arguments

Functions with a variable number of arguments present a particular portability problem if the type of the arguments is variable too. In such cases

the code is dependent upon the size of various data types.

In XENIX there is an include file, */usr/include/varargs.h*, that contains macros for use in variable argument functions to access the arguments in a portable way:

```
typedef char *va_list;
#define va_dcl int va_alist;
#define va_start(list) list = (char *) &va_alist
#define va_end(list)
#define va_arg(list,mode) ((mode *)(list += sizeof(mode)))[-1]
```

The va_end() macro is not currently required. Use of the other macros will be demonstrated by an example of the *fprintf* library routine. This has a first argument of type **FILE** *, and a second argument of type **char** *. Subsequent arguments are of unknown type and number at compilation time. They are determined at run time by the contents of the control string, argument 2.

The first few lines of *fprintf* to declare the arguments and find the output file and control string address could be:

```
#include <varargs.h>
#include <stdio.h>

int
fprintf(va_alist)
va_dcl;
{
        va_list ap;         /* pointer to arg list        */
        char *format;
        FILE *fp;

        va_start(ap);       /* initialize arg pointer */
        fp = va_arg(ap, (FILE *));
        format = va_arg(ap, (char *));

                ...
}
```

Note that there is just one argument declared to *fprintf*. This argument is declared by the va_dcl macro to be type **int**, although its actual type is unknown at compile time. The argument pointer "ap" is initialized by *va_start* to the address of the first argument. Successive arguments can be picked from the stack so long as their type is known using the *va_arg* macro. This has a type as its second argument, and this controls what data is removed from the stack, and how far the argument pointer "ap" is incremented. In *fprintf*, once the control string is found, the type of subsequent arguments is known and they can be accessed sequentially by repeated calls to va_arg(). For example, arguments of type **double**, **int** *, and **short**, could be retrieved as follows:

```
double dint;
int *ip;
short s;

dint = va_arg(ap, double);
ip = va_arg(ap, (int *));
s = va_arg(ap, short);
```

The use of these macros makes the code more portable, although it does assume a certain standard method of passing arguments on the stack. In particular no holes must be left by the compiler, and types smaller than **int** (e.g., **char**, and **short** on long word machines) must be declared as **int**.

B.4.7 Side Effects, Evaluation Order

The C language makes few guarantees about the order of evaluation of operands in an expression, or arguments to a function call. Thus

```
func(i++, i++);
```

is extremely nonportable, and even

```
func(i++);
```

is unwise if *func* is ever likely to be replaced by a macro, since the macro may use "i" more than once. There are certain XENIX macros commonly used in user programs; these are all guaranteed to use their argument once, and so can safely be called with a side-effect argument. The most common examples are *getc*, *putc*, *getchar*, and *putchar*.

Operands to the following operators are guaranteed to be evaluated left to right:

```
,       &&      ||      ?       :
```

Note that the comma operator here is a separator for two C statements. A list of items separated by commas in a declaration list is not guaranteed to be processed left to right. Thus the declaration

```
register int a, b, c, d;
```

on a PDP-11 where only three register variables may be declared could make any three of the four variables register type, depending on the compiler. The correct declaration is to decide the order of importance of the variables being register type, and then use separate declaration statements, since the order of processing of individual declaration statements is guaranteed to be sequential:

```
register int a;
register int b;
register int c;
register int d;
```

B.5 Program Environment Differences

Most programs make system calls and use library routines for various services. This section indicates some of those routines that are not always portable, and those that particularly aid portability.

We are concerned here primarily with portability under the XENIX operating system. Many of the XENIX system calls are specific to that particular operating system environment and are not present on all other operating system implementations of C. Examples of this are *getpwent* for accessing entries in the XENIX password file, and *getenv* which is specific to the XENIX concept of a process' environment.

Any program containing hard-coded pathnames to files or directories, or user IDs, login names, terminal lines or other system dependent parameters is nonportable. These types of constant should be in header files, passed as command line arguments, obtained from the environment, or obtained by using the XENIX default parameter library routines *dfopen*, and *dfread*.

Within XENIX, most system calls and library routines are portable across different implementations and XENIX releases. However, a few routines have changed in their user interface. The XENIX library routines are usually portable among XENIX systems.

Note that the members of the printf family, *printf*, *fprintf*, *sprintf*, *sscanf*, and *scanf* have changed in several ways during the evolution of XENIX, and some features are not completely portable. The return values of these routines cannot be relied upon to have the same meaning on all systems. Some of the format conversion characters have changed their meanings, in particular those relating to uppercase and lowercase in the output of hexadecimal numbers, and the specification of **long** integers on 16-bit word machines. The reference manual page for *printf* contains the correct specification for these routines.

B.6 Portability of Data

Data files are almost always nonportable across different machine CPU architectures. As mentioned above, structures, unions, and arrays have varying internal layout and padding requirements on different machines. In addition, byte ordering within words and actual word length may differ.

The only way achieve data file portability is to write and read data files as one dimensional character arrays. This avoids alignment and padding problems if

the data is written and read as characters, and interpreted that way. Thus ASCII text files can usually be moved between different machine types without too many problems.

B.7 Lint

Lint is a C program checker which attempts to detect features of a collection of C source files that are nonportable or even incorrect C. One particular advantage of *lint* over any compiler checking is that *lint* checks function declaration and usage across source files. Neither compiler nor loader do this.

Lint will generate warning messages about nonportable pointer arithmetic, assignments, and type conversions. Passage unscathed through *lint* is not a guarantee that a program is completely portable.

B.8 Byte Ordering Summary

The following conventions are used in the tables below:

a0 The lowest physically addressed byte of the data item. a0 + 1, and so on.

b0 The least significant byte of the data item, 'b1' being the next least significant, and so on.

Note that any program that actually makes use of the following information is guaranteed to be nonportable!

Byte Ordering for Short Types

CPU	Byte Order	
	a0	a1
PDP-11	b0	b1
VAX-11	b0	b1
8086	b0	b1
286	b0	b1
M68000	b1	b0
Z8000	b1	b0

Byte Ordering for Long Types

CPU	Byte Order			
	a0	a1	a2	a3
PDP-11	b2	b3	b0	b1
VAX-11	b0	b1	b2	b3
8086	b0	b1	b2	b3
286	b0	b1	b2	b3
M68000	b3	b2	b1	b0
Z8000	b3	b2	b1	b0

Appendix C
Building a Communication System

C.1 Introduction C-1

C.2 What You Need C-1

C.3 Installing the Modem C-2
- C.3.1 Choose a Serial Line C-2
- C.3.2 Set the Dialing Configuration C-3
- C.3.3 Connect the Modem C-3
- C.3.4 Test the Modem C-3

C.4 Creating a Dial-in Line C-5

C.5 Creating a Dial-out Line C-5
- C.5.1 Create the Call Unit Files C-6
- C.5.2 Create the L-devices File C-7
- C.5.3 Enable the Serial Line C-7

C.6 Installing a Uucp System C-8
- C.6.1 Choose a Uucp Site Name C-8
- C.6.2 Create the systemid File C-9
- C.6.3 Create a Dial-in Site C-9
- C.6.4 Create a Dial-Out Site C-11
- C.6.5 Linking Micnet Sites C-16

C.7 Maintaining the System C-17
- C.7.1 Displaying and Merging Log Files C-17
- C.7.2 Cleaning the Uucp Spool Directory C-18
- C.7.3 Reclaiming Log Files After a Crash C-18
- C.7.4 Reclaiming Data Files After a Crash C-19
- C.7.5 Checking the Transmission Status C-19
- C.7.6 Checking For Locked Sites or Devices C-20
- C.7.7 Creating Maintenance Shell Files C-20

C.8 Details of Operation C-20
C.8.1 Uucp Programs C-21
C.8.2 Uucp Directories and Files C-21
C.8.3 Uucp – Site to Site File Copy C-22
C.8.4 Uux – Site To Site Execution C-24
C.8.5 Uucico – Copy In, Copy Out C-26
C.8.6 Uuxqt – Uucp Command Execution C-30
C.8.7 Security C-30

C.9 Creating a New dial Program C-30

C.1 Introduction

This appendix explains how to build a communication system for your computer using a normal telephone line and a Hayes Smartmodem 1200. A communication system provides a way to:

- Log in to the computer from a remote terminal or computer.
- Use the **cu** command to call and log in to other computers.
- Use the **uucp** command to copy files to and from remote computers.
- Use the **uux** command to execute a remote mail program (**rmail**) on a remote computer.

In other words, the communication system is intended to give access to terminals and computers that cannot be connected to your computer through a direct serial line. In particular, the communication system is a practical solution to the problem of two Micnet networks (see the XENIX *Operations Guide*) that cannot be connected because of distance or cost of cable.

All communication tasks are supported by a variety of files and directories. In addition, the tasks invoked by the **uucp** and **uux** commands are actually performed by a system of underlying programs, called the uucp system. The files and underlying programs are described in full later in this appendix.

The following sections explain how to install the modem and prepare the programs you need to build a communication system. They also explain how to install and maintain a uucp system. The last section of this appendix presents the program listing of the **dial** program used to communicate with the modem. This listing may be used to create a modified **dial** program that communicates at a different baud rate (e.g., 300 baud) or with a different modem.

C.2 What You Need

To install a communication system on your computer, you will need

— A Hayes Smartmodem 1200

— A standard telephone jack for access to the telephone system (touch tone line required)

— An RS-232 serial line (or serial port) on your computer

— An RS-232 cable to connect the serial line to the modem

For proper operation of the modem, the RS-232 cable must provide the pin connections shown below. Note that the computer's serial connector must have

a DTE (Data Terminal Equipment) configuration. The modem is assumed to have a DCE (Data Communications Equipment) configuration.

Pin Connections

Computer (DTE)	Modem (DCE)
1	1
2	2
3	3
6	6
7	7
8	8
20	20

These pin connections are explained in the *Hayes Smartmodem 1200 Reference Manual*.

Make sure you inform the telephone company of your intent to use a modem with your telephone line. See the Hayes Reference Manual for details.

Finally, since many of the tasks you must perform require special permissions, you must log in to your computer's super-user account before performing them. Check with your computer's system manager before proceeding with this installation, or turn to the XENIX *Operations Guide* for instructions on how to log in as the super-user.

C.3 Installing the Modem

Installing the modem is the first step in creating a complete communication system. The installation has four steps:

1. Choose a serial line.
2. Set the dialing configuration.
3. Connect the modem.
4. Test the connection.

The following sections explain each step in detail.

C.3.1 Choose a Serial Line

You must choose the RS-232 serial line you wish to use with the system and connect to the modem. If there are no lines available, you must install a new serial line or make one available by removing any device connected to it. If you

remove a terminal, make sure no one is logged in.

Once you have chosen a serial line, find the name of the device special file associated with the line by looking in Appendix A of the XENIX *Operations Guide*. The filename should have the form

/dev/tty*nn*

where *nn* is the number of the corresponding line. For example, */dev/tty00* usually corresponds to serial line 0. You need the name for later steps.

C.3.2 Set the Dialing Configuration

In this communication system, your modem can be used to both send and receive calls. You must set the appropriate switches on the modem. Follow these steps:

1. Remove the front cover of the modem and locate the 8-pin configuration switch. (See the Hayes Reference Manual for instructions on how to remove the cover and the locate the switch.)

2. Set the pins on the configuration switch to the following positions:

pins	1	2	3	4	5	6	7	8
positions	up	up	down	down	up	up	up	down

3. Replace the front cover.

C.3.3 Connect the Modem

Once your modem's dialing configuration is set, you are ready to connect the modem to your computer. Review the installation instructions given in the Hayes Reference Manual, then follow these steps:

1. Connect the RS-232 serial cable to the serial line connector on the modem, then to the serial line connector on your computer. Make sure the cable is fully connected.

2. Plug the telephone line cable into the telephone connector on the modem, then into the telephone wall jack.

3. Plug in the power cord of the modem.

C.3.4 Test the Modem

As the last step of the modem installation, you should test the modem to make

sure that it can send and receive calls. Once you have verified that the modem is working, you can begin to use the communication system.

To test the modem, follow these steps:

1. Start the computer and log in as the super-user.

2. Disable the modem's serial line by typing

   disable /dev/tty*nn*

   where *nn* is the serial line's number.

3. Turn on the modem's power.

4. Make sure the volume switch on the modem is at an appropriate level. You must be able to hear the modem to carry out this test successfully. See the Hayes Reference Manual for the location of this switch.

5. Invoke the **dial** program using a command line of the form

   /usr/lib/uucp/dial /dev/tty*nn* *number speed*

   where /dev/tty*nn* is the filename of your serial line, and *number* is your telephone number (the number of the telephone jack your modem is connected to). For example, if your serial line is */dev/tty00* and number is "5551234", type

   /usr/lib/uucp/dial /dev/tty00 5551234 1200

6. Listen carefully to the modem. You should hear each digit as the number is dialed, then hear the busy signal when the telephone system tries to make connection with your modem.

7. If the busy signal is present, wait a few moments and listen carefully for the modem to hang up. The modem automatically discontinues any call it cannot make a connection for.

8. If the busy signal is not present, make sure you have connected the modem to the telephone jack. Make sure the jack is connected to the phone system. Make sure you gave the correct number when invoking **dial**.

9. If you did not hear the modem dial, make sure the volume switch is up. Make sure the modem is connected to the correct serial line and that the cable connection is tight. Make sure you gave the correct filename when invoking **dial**. Make sure modem's power is on.

C.4 Creating a Dial-in Line

You can create a dial-in line for use by remote terminals or computers by enabling the modem's serial line with the **enable** command after making sure your modem's serial line has an appropriate */etc/ttys* file entry. Once the line is enabled, any user at a remote terminal or computer can log in to your computer by calling your modem and following the ordinary login procedure. To create a dial-in line, follow these steps:

1. Log in as the super-user.
2. Use the **cat** command to examine the contents of the */etc/ttys* file.
3. Locate the entry in this file that corresponds to your serial line. The correct entry contains the name of your serial line and must have the form

   03tty*nn*

   where *nn* is the number of your serial line.
4. If necessary, use a XENIX text editor to change the entry or create a new entry for your serial line. Make sure the first two digits in the entry are "0" and "3", respectively.
5. Save the edited file and exit the editor.
6. Type

   enable /dev/tty*nn*

   where *nn* is the number of your serial line. This enables the line for logins.

Your computer will now receive calls from remote terminals or computers and prompt for a login name.

C.5 Creating a Dial-out Line

You can create a dial-out line by disabling the modem's serial line, creating the call unit files, and finally creating the */usr/lib/uucp/L-devices* file. A dial-out line lets you call and log in to other computers by using the **cu** command. The **cu** command uses the */usr/lib/uucp/L-devices* file to locate the modem's serial line and set the proper line speed when these values are not explicitly given in the **cu** command line. The call unit files are actually specially named files that are linked to your modem's serial line and used by uucp system programs.

The following sections explain how to create the necessary files and enable your line.

Note

Your modem's serial line cannot be both dial-in and dial-out at the same time. However, you can alternate between dial-in and dial-out at different times of the day by enabling or disabling the serial line as needed. Make sure you wait at least one minute between each invocation of the **enable** and **disable** commands.

C.5.1 Create the Call Unit Files

You must create two new device files, called "call unit" files, using the **ln** command and your serial line file. Follow these steps:

1. Log in as the super-user.

2. Check for any existing call unit files by using the **l** command. Type

   ```
   l /dev/cu*
   ```

   and examine the output. Call unit filenames have the form

   ```
   /dev/cua*n*
   ```

   and

   ```
   /dev/cul*n*
   ```

   where *n* is the same number as the corresponding serial line. If these files exist you can skip the next step and continue with creating the **/usr/lib/uucp/L-devices** file.

3. Use the **ln** command to create the call unit files. For example, if your serial line is named */dev/tty00*, type the commands

   ```
   ln /dev/tty00 /dev/cua0
   ```

   and

   ```
   ln /dev/tty00 /dev/cul0
   ```

4. Use the **chmod** command to change the access mode of the call unit files to read and write for everyone. For example, the command

```
chmod ugo+rw /dev/cua0 /dev/cul0
```

sets the appropriate permissions for the */dev/cua0* and */dev/cul0* files.

C.5.2 Create the L-devices File

The */usr/lib/uucp/L-devices* file defines the devices you intend to use to implement the dial-out line. The file is also used by programs in the uucp system (as described later). The file contains one or more entries of the form

type line call-unit speed

where *type* must be "ACU" if you are using an automatic call unit (modem) or "DIR" if you are using a direct serial line, *line* and *call-unit* are the */dev/cul* and */dev/cua* call unit filenames, respectively, and *speed* is the line speed or baud rate for transmissions. The call unit files are assumed to be in the */dev* directory, so the full pathname is not required. For example, the entry

```
ACU  cul0  cua0  1200
```

defines the device "/dev/cul0" as the line, "/dev/cua0" as the call-unit, and "1200" as the line speed.

With the Hayes modem, the *speed* must be set to 1200. Note that if you adapt the communication system for other purposes and give a *line* to a hardwired device (e.g., a direct serial line to another computer), you must use the number "0" for the *call-unit* field instead of a device name.

Use a XENIX text editor to create the file. Make sure you create the file in the */usr/lib/uucp* directory. Then use the **chmod** command to give the file read permission for everyone.

C.5.3 Enable the Serial Line

Enabling the serial line is the last step in creating a dial-out line. To enable the serial line, follow these steps:

1. Make sure your modem has been installed and tested.
2. Make sure you are logged in as the super-user.
3. Disable the modem's serial line by typing

   ```
   disable /dev/ttynn
   ```

   where *nn* is the number of your modem's serial line. If the line was already disabled, the command displays an error message that you

can safely ignore.

You will now be able to call other computers that have dial-in lines by using the **cu** command. For a complete description of the command, see *cu*(C) is the XENIX *Reference Manual.*

C.6 Installing a Uucp System

A uucp system is a set of files and programs that let you use the **uucp** and **uux** commands to transfer files and commands between computers connected by your communication system. Before you can use the **uucp** and **uux** commands, you must install the uucp system by creating or modifying a number of uucp system files.

The uucp system actually provides two different methods of interaction with other computers. One method requires a dial-in line through which remote computers can log in and transfer files and commands. With this method, your computer is called a "dial-in site." The other method requires a dial-out line through which your computer can call other computers. With this method, your computer is called a "dial-out site." Each method requires its own set of uucp system files.

Although you can install files for both methods, only one method can be used at a time. It is possible, however, to use alternate methods at different times of the day by creating a shell script that automatically enables or disables the line to permit dialing in or dialing out.

The following sections explain how to create files for both methods of interaction. They also explain how to create a transmission schedule and development a **cron** script to implement the schedule.

C.6.1 Choose a Uucp Site Name

In a uucp system, every computer belongs to a given "site". A site is any computer or any Micnet network that can communicate with the uucp system through a modem. To distinguish one site from another, every site must have a unique "site name". A site name is any combination of letters and digits that begins with a letter and is no more than seven characters long. The site name may then be used in **uucp** and **uux** commands to direct transmissions to the appropriate computer or Micnet network.

The site name should suggest some characteristic of the site, such as its location or affiliation. For example, a site in Chicago can be named "chicago", or a site in the shipping department can be named "shipping". The site name must be unique. That is, no other computer that calls your computer or is called by your computer can have the same site name.

Once you have chosen a site name, you will need to add it to the */etc/systemid* file as described in the next section.

C.6.2 Create the systemid File

Each site must have a */etc/systemid* file. The file defines the site name of the given site and associates the site with a Micnet network if any. The file has the form

sitename
[*machinename*]

where *sitename* is the name of the given site, and *machinename* is the Micnet machine name for that computer. The machine name is optional only if the computer is not connected to a Micnet network. For example, the entries

chicago
brewster

define a site named "chicago" whose Micnet machine name is "brewster".

Since uucp systems are often created after a Micnet network has been established, the *systemid* file usually already exists on a given site. In this case, you must add the site name to the beginning of each *systemid* file on each computer in the Micnet network. You may use a XENIX text editor. Note that you may give more than one machine name if desired, but each name must be on a separate line. For a full description of the *systemid* file, see *systemid*(M) in the XENIX *Reference Manual*.

C.6.3 Create a Dial-in Site

You can create a dial-in site by installing the uucp login information required by other computers that wish to log in and transfer files and commands. This information consists of the following:

- One or more */etc/passwd* file entries.
- User access information in the */usr/lib/uucp/USERFILE* file.

You can create this information by using a XENIX text editor and modifying or creating the appropriate files. The following sections explain the required format of the information.

Once the information is installed, you can enable the system for logins by creating a dial-in line. See the section "Creating a Dial-In Line" given earlier in this appendix.

Create Uucp Login Entries

A dial-in site must provide a login entry for the sites that call it. These entries must be placed in the */etc/passwd* file. A uucp login entry has the same form as an ordinary user login entry (see Chapter 3 in the XENIX *Operations Guide*), but gives a special login directory and login program instead of the normal user directory and shell. To create a uucp login entry, follow these steps:

1. Choose a new login name and a user ID for the uucp login. The name may be any combination of letters and digits that is no more than eight characters long. The user ID must be an integer number in the range 1 to 65535. Make sure the name and ID are unique; a uucp login entry must not have the same name or ID as any other login entry.

2. Invoke a XENIX text editor giving */etc/passwd* as the file to edit.

3. Move to the end of the file and insert the login entry using the form

   login-name::*user-ID*:*group-ID*::/usr/spool/uucp:/usr/lib/uucp/uucico

   where *login-name* is the login name you have chosen, and *user-ID* and *group-ID* are both the user ID you have chosen. For example, if you have chosen "uuchcg" for the login name and "12" for the user ID, add the entry

   uuchcg::12:12::/usr/spool/uucp:/usr/lib/uucp/uucico

   to the end of the file.

4. Save the new file and exit the editor.

5. Create a new password for the login with the **passwd** command. Type

   passwd *login-name*

   where *login-name* is the login name you have chosen. The command will ask you to type the new password twice. It will then add the encrypted password to the new login entry.

Note that you can create new login entries for each site that calls your site, or use one entry for all sites.

Create the USERFILE

The */usr/lib/uucp/USERFILE* file defines which directories a given site (or a given user) may access using the **uucp** and **uux** commands. You should create one *USERFILE* entry for each site or user with a login entry in the */etc/passwd*

file. Each entry has the form

login,sitename pathname...

where *login* is the login name for a given site, *sitename* is the site name of a given site, and *pathname* is the full pathname of the directory the given site may access. More than one pathname may be given if desired. The *login* and *sitename* are optional.

The following rules explain how access is granted for each entry.

1. A calling site is granted access to those directories defined in an entry containing its site name.

2. A calling site whose name does not appear in an entry is granted access to the directories defined for the first entry with no site name.

3. A user is granted access to those directories defined in an entry containing his login name.

4. A user whose login name does not appear in an entry is granted access to directories defined in the first entry with no login name.

You may have more than one entry with the same login name if you wish, but you must make sure that at least one of these entries also has the site name of any calling site which can log in with that name, or that one of these entries has no site name.

For example, consider the following entries.

```
uuccg,chicago  /usr  /usr2/market
uucp,      /usr/vendor
schmidt,  /usr2/market  /usr/vendor
,              /usr/spool/uucp/public
```

The site named "chicago" has access to files in the directories named "/usr" and "/usr2/market". Any other site will be granted access to "/usr/vendor" only. A local user named "schmidt" is granted access to the directories "/usr2/market" and "/usr/vendor". All other users have access to "/usr/spool/uucp/public" only.

C.6.4 Create a Dial-Out Site

You can create a dial-out site by installing the dialing information needed by your system to call and log in to other computers. This information consists of the following:

- Dialing abbreviations for remote computers in the */usr/lib/uucp/L-dialcodes* file.

- Information about logins on remote computers in the */usr/lib/uucp/L.sys* file.
- A transmission schedule in the form of a shell script to be called periodically by the **cron** program.

You can create this information by using a XENIX text editor and modifying or creating the appropriate files. The following sections explain the required format of the information.

Once the information is installed, you can enable the system for calling other computers by creating a dial-out line. See the section "Creating a Dial-Out Line" given earlier in this appendix.

Create the L-dialcodes File

The */usr/lib/uucp/L-dialcodes* file defines abbreviations for often used telephone prefixes and area codes. You may use these abbreviations in the *L.sys* file when forming the telephone numbers of remote sites.

The *L-dialcodes* file may contain one or more entries of the form

abbreviation dial-sequence

where *abbreviation* is any combination of letters and digits that begins with a letter, and *dial-sequence* is any combination of digits that represents a telephone prefix, area code or any other part of a telephone number. For example, the entry

ms 555

defines the abbreviation "ms" to be the telephone prefix "555".

Create the L.sys File

The */usr/lib/uucp/L.sys* file defines the names, telephone numbers, and login information of all sites in the system. The file contains one or more entries of the form

sitename time device speed phone login

where *sitename* is the name of the site to be called, *time* is a combination of letters and digits that gives the weekdays and times when the given site can be called, *device* is the name of the device through which the given site is to be called, *speed* is the line speed for the call, *phone* is the phone number of the given site, and *login* is the login information required to log in to the given site. With the Hayes modem, the *speed* must be **1200**.

The *time* defines when the given site can make calls to other sites. It has the form

days times

where *days* is a list of one or more days of the week, and *times* is a range of times of day. The days of the week may be "Su", "Mo", "Tu", "We", "Th", "Fr", "Sa", "Wk" (for any week-day), "Any" (for any day), and "Never" (for call by special request only). The time of day must be given as a four digit number. The first pair of digits gives the hour (in terms of a 24 hour clock), the second pair gives the minutes. A range of times is a pair of times of the day separated by a hyphen (–). For example, the entry

MoTuTh0800–1230

allows the given site to be called any Monday, Tuesday, or Thursday from 8 in the morning to 12:30 in the afternoon.

The *device* must be the keyword "ACU" if you are using a modem. If you are using a direct line to the other site, then you must give the filename of the serial line (or other device) you intend to use (e.g. tty01).

The *phone* must be the telephone number of the given site. It must have the correct number of digits (including area code if necessary) or be a combination of *L-dialcodes* abbreviations and digits. *L-dialcode* abbreviations must go before any digits. Hyphens must not be used. For example, "5551234" is a valid local number and "2065551234" is a valid long distance number. If the abbreviation "ms" is defined to be "555", then "ms1234" may be used in place of "5551234".

With the Hayes modem, you may use a comma (,) in a number to cause a delay when dialing. This is useful if you must dial for an outside line before placing the call. For example, the number "9,5551234" causes a delay immediately after the "9" has be dialed. After the delay, the rest of the number is dialed. If you are not using a modem, then *phone* must be the filename of the device you intend to use instead of a phone number.

The *login* must be a sequence of names, numbers, and other information that represents the steps required to log in to the given site. This sequence has the form

expect send [*expect send*]...

where *expect* is the prompt or message that you expect the given site to return to the calling site, and *send* is the name, number, or other information that you wish to send in response to the expected prompt or message. For example, the following is the login sequence for a typical XENIX site

login: uuccg ssword: market

Note that "ssword:" is given instead of the complete prompt "Password:". Only the last eight characters in each expected prompt or message are examined, so you do not need to give the preceding characters if you wish to save space.

If you anticipate problems during the login sequence, you may include a conditional response immediately after each expected prompt or message. This conditional response has the form

expect [*–send–expect1*] ...

where *expect* is the prompt or message you expect the given site to return, *send* is the name or number you wish to send if the prompt or message returned is not correct, and *expect1* is the prompt or message you expect after sending the conditional response. For example, the following shows how to invoke the "login" prompt if it is not immediately present.

--login-EOT-login-uuccg ssword: market

There are two special keywords that you may use in the login sequence. The "EOT" keyword causes an end of transmission character to be sent and the "BREAK" keyword causes a break character to be sent. (The break character is simulated using line speed changes and null characters and may not work on all devices and/or systems.)

The complete *L.sys* entry must be placed on one line as shown by the following example.

chicago Any ACU 1200 5551234 login uucp ssword: market

Create a Transmission Schedule

In the uucp system, the **uucico** program carries out all transmissions between your site and other sites, sending and receiving files and commands as long as there is work for it to do. On a dial-in site, **uucico** is always started whenever a calling site logs in, but on a dial-out site, **uucico** is only started when an explicit invocation of the program is given. This means you must periodically invoke the program on a dial-out site to ensure that all transmissions requested by the **uucp** and **uux** programs are completed. You can do this in one of two ways: invoke the program manually whenever you need it, or create a shell script and let the **cron** program invoke **uucico** automatically according to a schedule of transmissions.

The most convenient method is to let **cron** invoke **uucico** for you. To do this, you must choose a schedule of times for **uucico** to be invoked, then create a */etc/crontab* file entry for this schedule. A */etc/crontab* entry has the form

minutes hour day month day-of-week command-line

where *minutes*, *hour*, *day*, *month*, and *day-of-week* give the exact day of the year and time of day to execute the given *command-line*. Each item, except the *command-line*, must be an integer number within an acceptable range, e.g., 0 to 59 for *minutes*. A sequence of values for one item may be given by separating the values with commas. Also, an asterisk (*) may be given to represent all acceptable values. The *command-line* must be the name of the shell script you have created to invoked **uucico**.

You can add an entry to the */etc/crontab* file by using a XENIX text editor. For more information about the file, see *cron*(C) in the XENIX *Reference Manual*. For example, the entry

15,45 * * * * /usr/lib/uucp/transmit

invokes the shell script "transmit" every 30 minutes to sites for which requests are pending. The entry

0 0 * * * /usr/lib/uucp/transmit

invokes "transmit" every day at midnight, and the entry

15 2,4,6 * * * /usr/lib/uucp/transmit

invokes the script every day at "2:15", "4:15", and "6:15" in the morning.

A shell script is simply a text file that contains one or more XENIX commands. If your uucp system is acting as both a dial-in and dial-out site at different times, then the script should have the form

```
disable /dev/ttynn
uucico -r1 -ssitename
        .
        .
        .
enable /dev/ttynn
```

where *nn* is the number of your modem's serial line, *sitename* is the name of the site you wish to call. Use the –s option if you wish to force a call to the given site even if no requests for transmissions exist on the calling site. Note that the –S option may be used in place of the –s option if you wish to ignore the range of calling times given in the *L.sys* file. Give one **uucico** command for each site you wish to call. If you want to call only those sites for which requests exist, give a single **uucico** command, but do not give the –s or –S option with the command. If your computer is strictly a dial-out site, then the **enable** and **disable** commands are not required.

For example, the script

```
disable /dev/tty00
uucico -r1 -schicago
enable /dev/tty00
```

will place a call to the "chicago" site after disabling the serial line. The line must be disabled in order to dial out on that line. It must be enabled to allow subsequent calls from other computers.

You can create a shell script by using a XENIX text editor. For convenience, the script should be placed in the */usr/lib/uucp* directory and must be given execute permissions for everyone. Note that you can also add uucp maintenance programs to the script. See the section "Creating Maintenance Shell Files" later in this appendix.

C.6.5 Linking Micnet Sites

You can send and receive mail from other Micnet sites through the uucp system by defining a uucp alias in the *maliases* file of each computer in your Micnet site. A uucp alias is any alias having the form

sitename!

where *sitename* is the name of a Micnet site. Uucp aliases can be used in **mail** commands to direct mail through the uucp system to the desired Micnet site. See the XENIX *Operations Guide* for details.

To use a uucp system with your Micnet network, follow these steps:

1. Add the entry

   uucp:

   to the *maliases* file of the computer on which the uucp system is installed.

2. For all other computers in your site, add the entry

   uucp:*machine-name*?

   to the *maliases* file. The *machine-name* must be the name of the computer on which the uucp system is installed.

You can test the uucp alias by mailing a short letter to yourself via another site. For example, if you are on the site "chicago", and there is another Micnet site named "seattle" in the system, then the command

```
mail seattle!chicago!johnd
```

will send mail to the "seattle" site, then back to your "chicago" site, and finally to the user "johnd" in your Micnet network. Note that a uucp system usually performs its communication tasks according to a fixed schedule, and may not return mail immediately. If there are problems, check that the uucp installation is correct and that you have added the correct uucp aliases to all *maliases* files in your site. Also, make sure that the remote site has the correct uucp aliases in its *maliases* files.

C.7 Maintaining the System

This section explains how to maintain the uucp system. In particular, it explains how to display and merge the content of uucp log files, how to remove old requests and files from the spool directories, and how to solve some common problems.

You can automate some maintenance tasks by creating shell command files and initiating these files with *crontab* entries. Other tasks require manual modification. Some sample shell files are given toward the end of this section.

C.7.1 Displaying and Merging Log Files

You can display a record of the transmissions requested and completed to a give site or user by using the **uulog** command. The command displays the contents of the individual log files created for a given site or user and merges these entries with the system log file *LOGFILE*. The log files contain information about queued requests, calls to remote sites, execution of **uux** commands, and file copy results. The command has the form

uulog –s*sitename* –u*user*

where –s*sitename* names the site whose log files are to be displayed, and –u*user* names the user whose log files are to be displayed. If you do not give a *sitename* and *user* log files for all sites and users are displayed. The command places the new log files at the beginning of the existing *LOGFILE*.

The log files are originally created in the */usr/spool/uucp* directory as individual files, but should be copied to the *LOGFILE* on a regular basis since they are not copied automatically. For example, the command

uulog

merges all log files and displays their contents. The command

uulog –schicago

merges only log files created for the site "chicago".

Note that the system *LOGFILE* should be removed periodically since it is copied each time new log files are put into the file.

C.7.2 Cleaning the Uucp Spool Directory

You can remove unwanted uucp system files from the uucp spool directory by using the **uuclean** command. The command removes temporary data, *LOG*, system status, and lock files from the spool directory if they are more than a given number of hours old. The command has the form

uuclean –d*dir* –m –n*hours* –p*pre* –x*n*

where –d*dir* names the directory to be scanned, –*m* causes mail to be sent to the owner of each file removed, –n*hours* gives the age in hours of files to be removed, –p*pre* causes files with the given prefix to be examined and removed, and –x*n* directs the command to give the *n*th level of debugging output. Up to 10 file prefixes may be specified with the –p option. If –m is given most mail will be sent to the owner of the uucp programs since most files put into the spool directory will be owned by the owner of the uucp programs. This is a result of the setuid bit being set on these programs. The default number of *hours* is 72 (3 days).

The **uuclean** program should be run once a day. You can invoke it automatically by using a system daemon such as **cron**. The command

uuclean –pTM

removes all temporary data files that are at least three days old. The command

uuclean –pLCK –h1 –m

removes all lock files that are at least an hour old and mails a list of each file removed to the owner.

The **uuclean** command may also be run as needed to remove unwanted files after a system crash or an aborted uucp program.

C.7.3 Reclaiming Log Files After a Crash

You can reclaim individual log files after a system crash by changing their access mode with the **chmod** command, then using **uulog** command. After a transmission failure or system crash, the individual log file for the transmission may be left with access mode 0222 making it impossible for the **uulog** command to read the file. To reclaim the log file you must, use **chmod** to change the access mode to 0666. You can then let **uulog** merge them with the *LOGFILE*.

C.7.4 Reclaiming Data Files After a Crash

You can check the status of files transmitted from a remote site and possibly reclaim some or all of the data lost during an aborted transmission by examining system data files. The data files contain the contents of files copied from remote sites. These files are temporarily kept in the */usr/spool/uucp* directory and their names have the form

TM.*pid.ddd*

where *pid* is a process-id and *ddd* is a sequential three digit number starting at zero for each invocation of **uucico** and incremented for each file received.

The temporary data files are normally moved to the requested destination immediately after the transmission has finished. However, if a transmission has failed or the system has crashed, the file remains in the spool directory. You can examine the contents of this file with the **cat** command. If desired, you can reclaim the file by moving it to a new location with the **mv** command. Leftover data files that cannot be reclaimed should be removed using the **uuclean** command.

C.7.5 Checking the Transmission Status

You can check the status of transmissions between sites in the uucp system by examining the system status files. System status files contain information about login, dialup, or sequence check failure, as well as the talking status when two machines are conversing. The files are kept in the */usr/spool/uucp* directory and their names have the form

STST.*sitename*

where *sitename* is the name of the remote site.

Normally, system status files are removed after each successful transmission, but when a failure occurs, the uucp system copies information about the failure to the file and leaves it in the directory. This prevents the uucp system from making further calls to the given site for about an hour, or for sequence check failures, until the file is removed.

To examine the status, use the **cat** command to display the contents of the file. If problems with transmissions are detected it may indicate a problem with the modem or with the serial line connected to the modem.

If a system status file has been left due to a program or system crash, the file may prevent all subsequent transmissions to the given site. In this case, the file must be removed before attempting further calls.

C.7.6 Checking For Locked Sites or Devices

You can make sure the uucp system is not intentionally preventing transmissions to a given site or through a given device by examining the system lock files. The uucp system creates a lock file for each site being called and for each device being used to call a site. Lock files prevent the uucp system from attempting to duplicate conversations with a given site, or from placing multiple calls on the same device. The lock files are kept in the */usr/spool/uucp* directory and their names have the form

LCK..*str*

where *str* is either a site name or the name of the calling device.

Since lock files prevent all calls to a given site or through a given device, it is wise to make sure no unnecessary lock files are left in the directory. If a transmission has been aborted or the system has crashed, the lock files will prevent subsequent transmissions for about about 24 hours. If you wish to place a call before this time you must remove the file using the **uuclean** command.

C.7.7 Creating Maintenance Shell Files

The **uulog** and **uuclean** command can be invoked automatically by placing the commands in a shell file and creating a *crontab* entry for the shell file. The system daemon **cron** will then invoke the commands at the given times and most of the simple maintenance will be performed. For example, you can create a shell file that daily removes *TM*, *ST*, and *LCK* files.
and *C.* or *D.* files for work which can not be accomplished for reasons such as bad phone number and login changes. In this case, the shell file should contain the commands

```
/usr/lib/uucp/uuclean -pTM -pC. -pD.
/usr/lib/uucp/uuclean -pST -pLCK -n12
```

Note that the **-n12** option causes the *ST* and *LCK* files older than 12 hours to be deleted. An appropriate *crontab* entry must be created in order to invoke the shell file automatically.

C.8 Details of Operation

This section describes the details of uucp system program operation. It explains the processes used to create system communication and defines the files used to support the system.

C.8.1 Uucp Programs

The uucp system consists of four primary and two secondary programs. The primary programs are

uucp	This program creates work and gathers data files in the spool directory for the transmission of files.
uux	This program creates work and execute files, and gathers data files for the remote execution of XENIX commands.
uucico	This program executes the work files for data transmission.
uuxqt	This program executes XENIX commands found in execution files.

The secondary programs are

uulog	This program updates the log file with new entries and reports on the status of uucp requests.
uuclean	This program removes old files from the spool directory.
dial	This program directs the modem to dial a remote site.

C.8.2 Uucp Directories and Files

During execution of the uucp programs, the uucp system uses files from the following three directories:

/usr/lib/uucp
This is the directory used for uucp system files and all executable programs other than **uucp** and **uux**.

/usr/spool/uucp
This is the spool directory used during *uucp* execution.

/usr/spool/uucp/.XQTDIR
This directory is used during execution of execute files.

Files are created in a spool directory for processing by the uucp daemons. There are three types of files used for the execution of work:

Data files	Contain data for transfer to remote sites
Work files	Contain directions for file transfers between sites
Execution files	Contain directions for XENIX command executions which involve the resources of one or more sites.

C.8.3 Uucp – Site to Site File Copy

The uucp program is the user's primary interface with the system. The uucp program was designed to look like the cp command. The syntax is

uucp [*option*] ... *source* ... *destination*

where *source* and *destination* may contain the prefix *sitename!* which indicates the site on which the file or files reside or where they will be copied.

The options interpreted by uucp are

-d Make directories when necessary for copying the file.

-c Don't copy source files to the spool directory, but use the specified source when the actual transfer takes place.

-m Send mail on completion of the work.

The following options are used primarily for debugging:

-s*dir* Use directory *dir* for the spool directory.

-x*num* Use *num* as the level of debugging output.

The destination may be a directory name, in which case the file name is taken from the last part of the source's name. The source name may contain special shell characters such as "?*[]". If a source argument has a *sitename!* prefix for a remote site, the file name expansion will be done on the remote site.

The command

uucp *.c chicago!/usr/dan

will set up the transfer of all files whose names end with *.c* to the */usr/dan* directory on the *chicago* machine.

The source and/or destination names may also contain a ~*user* prefix. This translates to the login directory on the specified site. For names with partial pathnames, the current directory is prepended to the file name. File names with "../" are not permitted.

The command

uucp chicago!~dan/*.h ~dan

will set up the transfer of files whose names end with *.h* in dan's login directory on site *usg* to dan's local login directory.

For each source file, the program will check the source and destination filenames and the site-part of each to classify the work into one of five types:

1. Copy source to destination on local site.
2. Receive files from other sites.
3. Send files to remote sites.
4. Send files from remote sites to another remote site.
5. Receive files from remote sites when the source contains special shell characters as mentioned above.

After the work has been set up in the spool directory, the **uucico** program must be started to try to contact the other machine to execute the work.

Copying Files to a Local Destination

A **cp** command is used to do type 1 work. The **–d** and the **–m** options are not honored in this case.

Receiving Files from Other Sites

For type 2 work, a one line work file is created for each file requested, and is put in the spool directory with the following fields, each separated by a blank:

[1] R

[2] The full pathname of the source or a ~user/pathname. The ~*user* part will be expanded on the remote site.

[3] The full pathname of the destination file. If the ~*user* notation is used, it will be immediately expanded to be the login directory for the user.

[4] The user's login name.

[5] A "–" followed by an option list. (Only the **–m** and **–d** options will appear in this list.)

Sending Files to Remote Sites

For type 3 work, a work file is created for each source file and the source file is copied into a data file in the spool directory. (A **–c** option on the **uucp** program will prevent the data file from being made. In this case, the file will be

transmitted from the indicated source.) Pathnames are checked using the *USERFILE* to verify access to the requested directory. The fields of each entry are given below.

[1] S

[2] The full pathname of the source file.

[3] The full pathname of the destination or ~user/filename.

[4] The user's login name.

[5] A "–" followed by an option list.

[6] The name of the data file in the spool directory.

[7] The file mode bits of the source file in octal print format (e.g. 0666).

Copying Files Between Sites

For type 4 and 5 work, **uucp** generates a **uucp** command line and sends it to the remote machine; the remote **uucico** executes the command line.

C.8.4 Uux – Site To Site Execution

The **uux** command is used to set up the execution of a XENIX command where the execution machine and/or some of the files are remote. The syntax of the **uux** command is

uux [–] [*option*] ... *command-string*

where *command-string* is made up of one or more arguments. All special shell characters such as "<>|^" must be quoted either by quoting the entire command string, or by quoting the character as a separate argument. Within the command string, the command and file names may contain a *sitename!* prefix. All arguments which do not contain a "!" will not be treated as files. (They will not be copied to the execution machine.) The – option is used to indicate that the standard input for the given command should be inherited from the standard input of the **uux** command. The only option is essentially for debugging: –x*num* directs the command to use *num* as the level of debugging output.

The command

pr abc | uux – chicago!rmail joe

will set up the output of "pr abc" as standard input to a **mail** command to be executed on site *usg*.

Uux generates an execute file which contains the names of the files required for execution (including standard input), the user's login name, the destination of the standard output, and the command to be executed. This file is either put in the spool directory for local execution or sent to the remote site using a generated send command (type 3 above).

For required files which are not on the execution machine, **uux** will generate receive command files (type 2 above). These command-files will be put on the execution machine and executed by the **uucico** program. (This will work only if the local site has permission to put files in the remote spool directory as controlled by the remote *USERFILE*.)

The execute file will be processed by the **uuxqt** program on the execution machine. It is made up of several lines, each of which contains an identification character and one or more arguments. The order of the lines in the file is not relevant and some of the lines may not be present. Each line is described below.

User Line

U *user site*

where the *user* and *site* are the requestor's login name and site.

Required File Line

F *filename real-name*

where the *filename* is the generated name of a file for the execute machine and *real-name* is the last part of the actual file name (contains no path information). Zero or more of these lines may be present in the execute file. The **uuxqt** program will check for the existence of all required files before the command is executed.

Standard Input Line

I *filename*

The standard input is either specified by a "<" in the command-string or inherited from the standard input of the **uux** command if the – option is used. If a standard input is not specified, */dev/null* is used.

Standard Output Line

O *filename sitename*

The standard output is specified by a ">" within the command-string. If a standard output is not specified, */dev/null* is used. (Note that the use of ">>" is not implemented.)

Command Line

C *command* [*arguments*] ...

The *arguments* are those specified in the command string. The standard input and standard output will not appear on this line. All required files will be moved to the execution directory (a subdirectory of the spool directory) and the XENIX command is executed using the shell. In addition, a shell PATH statement is prepended to the command line as specified in the **uuxqt** program.

After execution, the standard output is copied or set up to be sent to the proper place.

C.8.5 Uucico – Copy In, Copy Out

The **uucico** program will perform the following major functions:

— Scan the spool directory for work.

— Place a call to a remote site.

— Negotiate a line protocol to be used.

— Execute all requests from both sites.

— Log work requests and work completions.

Uucico may be started by a system daemon, by the user (this is usually for testing), or by a remote site. (The **uucico** program should be specified as the shell field in the */etc/passwd* file for the uucp logins.)

When started with the **–r1** option, the program is considered to be in MASTER mode. In this mode, a connection will be made to a remote site. If started by a remote site, the program is considered to be in SLAVE mode.

The MASTER mode will operate in one of two ways. If no site name is specified (the **–s** option not specified) the program will scan the spool directory for sites to call. If a site name is specified, that site will be called, and work will only be done for that site.

The **uucico** program must generally started directly by the user or by another program, such as a shell script invoked by **cron**. There are several options used for execution:

–r1 Start the program in MASTER mode. This is used when **uucico** is started by a program or **cron** shell.

–s*sitename* Do work only for site *sitename*. If **–s** is specified, a call to the specified site will be made even if there is no work for site

sitename in the spool directory, but will only call when times in the *L.sys* file permit it. This is useful for polling sites which do not have the hardware to initiate a connection.

–S*sitename* Do work only for site *sitename*. If –S is specified a call to the specified site will be made even if there is no work for the site in the site in the spool directory. Unlike –s, this option ignores the call times for the *sitename* given in the *L.sys* file.

The following options are used primarily for debugging:

–d*dir* Use directory *dir* for the spool directory.

–x*num* Use *num* as the level of debugging output.

The next part of this section will describe the major steps within the **uucico** program.

Scanning For Work

The names of the work related files in the spool directory have the format

type . sitename grade number

where *type* may be "C" for copy command file, "D" for data file, "X" for execute file, *sitename* is the remote site, *grade* is a character, and *number* is a four digit, padded sequence number.

The file

C.res45n0031

is a work file for a file transfer between the local machine and the "res45" machine.

The scan for work is done by looking through the spool directory for work files (files with prefix "C."). A list is made of all sites to be called. **Uucico** calls the site specified by the –s or –S option and process the corresponding work files.

Calling a Remote Site

The call is made using information from several files which reside in the uucp program directory. At the start of the call process, a lock is set to forbid multiple conversations between the same two sites. The lock filename has the form

LCK..*str*

where *str* is the device name. The file is in the */usr/spool/uucp* directory.

The site name is found in the *L.sys* file. The information contained for each site is

[1] Site name

[2] Times to call the site (days-of-week and times-of-day)

[3] Device or device type to be used for call

[4] line speed

[5] phone number if field [3] is "ACU" or the device name (same as field [3]) if not

[6] Login information (multiple fields)

The time field is checked against the present time to see if the call should be made.

The *phone number* may contain abbreviations (e.g. mh, py, boston) which get translated into dial sequences using the *L-dialcodes* file.

The *L-devices* file is scanned using device type and line speed fields from the *L.sys* file to find an available device for the call. The program will try all devices which satisfy these fields until the call is made or until no more devices can be tried. If a device is successfully opened, a lock file is created so that another copy of **uucico** will not try to use it. If the call is complete, the login information in the last field of *L.sys* is used to login.

The conversation between the two **uucico** programs begins with a handshake started by the SLAVE site. The SLAVE sends a message to let the MASTER know it is ready to receive the site identification and conversation sequence number. The response from the MASTER is verified by the SLAVE and if acceptable, protocol selection begins. The SLAVE can also reply with a call-back required message in which case, the current conversation is terminated.

Selecting Line Protocol

The remote site sends a message

P*proto-list*

where *proto-list* is a string of characters, each representing a line protocol.

The calling program checks the protocol list for a letter corresponding to an available line protocol and returns a use protocol message. The message has the form

U*code*

where *code* is either a one character protocol letter or "N" which means there is no common protocol.

Processing Work

The initial role of MASTER or SLAVE for the work processing is the mode in which each program starts. (The MASTER has been specified by the **–r1** option.) The MASTER program does a work search similar to the one used in the section "Scanning For Work" above.

There are five messages used during the work processing, each specified by the first character of the message. They are

S	Send a file
R	Receive a file
C	Copy complete
X	Execute a **uucp** command
H	Hangup

The MASTER will send "R", "S", or "X" messages until all work from the spool directory is complete, at which point an "H" message is sent. The SLAVE will reply with the first letter of the request and either the letter "Y" or "N" for yes or no. For example, the message "SY" indicates that it is okay to send a file.

The send and receive replies are based on permission to access the requested file/directory using the *USERFILE* and read/write permissions of the file/directory. After each file is copied into the spool directory of the receiving site, a copy-complete message is sent by the receiver of the file. The message "CY" will be sent if the file has successfully been moved from the temporary spool file to the actual destination. Otherwise, a "CN" message is sent. (In the case of "CN", the transferred file will be in the spool directory with a name beginning with "TM".) The requests and results are logged on both sites.

The hangup response is determined by the SLAVE program by a work scan of the spool directory. If work for the remote site exists in the SLAVE's spool directory, an "HN" message is sent and the programs switch roles. If no work exists, an "HY" response is sent.

Terminating a Conversation

When an "HY" message is received by the MASTER it is echoed back to the SLAVE and the protocols are turned off. Each program sends a final "OO"

message to the other. The original SLAVE program will clean up and terminate. The MASTER will proceed to call other sites and process work as long as possible or terminate if a –s option was specified.

C.8.6 Uuxqt – Uucp Command Execution

The **uuxqt** program is used to process execute files generated by **uux**. The **uuxqt** program is started by the **uucico** program. The program scans the spool directory for execute files (prefix *X.*). Each one is checked to see if all the required files are available and if so, the command line or send line is executed.

The execute file is described in the section "Uux - Site to Site Execution" above.

The execution is accomplished by executing the shell command

```
sh -c
```

with the command line after appropriate standard input and standard output have been opened. If a standard output is specified, the program will create a send command or copy the output file as appropriate.

C.8.7 Security

In unrestricted uucp system, once a user logs in to another site through the uucp system, the user can execute any commands and copy any files normally accessible to the uucp login. It is up to the individual sites to be aware of this and apply the protections that they feel are necessary to prevent unauthorized use of files and commands.

The uucp system does provide a certain level of security. For example, a calling site does not get a standard shell when it logs in. Instead, the **uucico** program is started and all work is done through this program. The **uucico** program checks the pathnames of file to be sent or received to prevent access to restricted directories. The *USERFILE* supplies the information for these checks. To prevent execution of possibly damaging commands, the **uuxqt** program can only execute the **rmail** program on a remote site. This special program is one of many underlying mail programs that help deliver mail. Finally, the *L.sys* file is owned by uucp and has mode **0400** to protect the phone numbers and login information for remote sites.

C.9 Creating a New dial Program

The **dial** program is used by dial-out sites to place calls to other computers. You can create a new **dial** program for any dial-out site not using the Hayes Smartmodem by modifying the source program given in Figure C.1 and compiling it with the **cc** command. See the XENIX *Programmer's Guide* for information about the **cc** command.

Figure C.1 – Source Program for dial

```
/*
 *
 *      Copyright (C) Microsoft Corporation, 1983
 *
 * Simple dialer program for the Hayes "Smart" Modem 1200
 *
 *      See Hayes manual for command definitions
 *
 * Usage:  dial ttyname telnumber speed
 *
 *      returns 0 if a connection was made
 *              -1 otherwise
 */

#include <stdio.h>
#include <signal.h>
#include <fcntl.h>
#include <sys/types.h>
#include <sys/ioctl.h>
#include <termio.h>

#define SAME    0
char *setup = "M1 F1 DT";               /* Speaker on, Full Duplex, Touch tone*/
struct termio term;
int baudrate;                           /* baud rate of modem */
char buffer[80];
int alrmint();

main(argc,argv)
int argc;
char *argv[];
{
        FILE *fdr,*fdw;
        int fd;

        if( argc != 4) {
                fprintf(stderr,"Usage: dial  devicename [number] speed\n");
                exit(-1);
        }
        if( (fd=open(argv[1],O_RDWR|O_NDELAY)) < 0 ) {
                fprintf(stderr,"dial: Can't open device: %s for reading.\n",argv[1]);
                exit(-1);
        }
        switch(atoi(argv[3])) {
                case 300:
                        baudrate = B300;
                        break;
                case 1200:
```

```
                        baudrate = B1200;
                        break;
                default:
                        baudrate = B1200;
        }
        /*
         * set line for no echo and specific speed
         */
        ioctl(fd, TCGETA, &term);
        term.c_cflag &= ~CBAUD;
        term.c_cflag |= CLOCAL|HUPCL|baudrate;
        term.c_lflag &= ~ECHO;
        term.c_cc[VMIN] = 1;
        term.c_cc[VTIME] = 0;
        ioctl(fd, TCSETA, &term);
        fcntl(fd, F_SETFL, fcntl(fd, F_GETFL, 0) & ~O_NDELAY);
        if( (fdr=fopen(argv[1],"r")) == (char *)NULL ) {
                fprintf(stderr,"dial: Can't open device: %s for reading.\n",argv[1]);
                exit(-1);
        }
        if( (fdw=fopen(argv[1],"w")) == (char *)NULL ) {
                fprintf(stderr,"dial: Can't open device: %s for writing.\n",argv[1]);
                exit(-1);
        }
        setbuf(fdw,0);                                  /* Want unbuffered I/O */
        /*
         * setup for timeout in 10 seconds if no response
         */
        signal(SIGALRM, alrmint);
        alarm(10);
reread:
        fprintf(fdw,"AT\r");                            /* Put Hayes into command mode
        if( fgets(buffer,sizeof buffer, fdr) == (char *)NULL )
                exit(-1);
        if( strncmp(buffer, "OK",2) != SAME ) {                 /* got back an OK? */
                sleep(1);
                goto reread;
        }
        alarm(0);                                       /* turn off alarm */
        sleep(1);
        fprintf(fdw,"AT %s %s\r",setup,argv[2]);        /* put out dialing string */
        /*
         * turn off CLOCAL now, since we want modem interrupts to work
         * setup alarm.  (Longer timeout period for longer numbers)
         *
        ioctl(fd, TCGETA, &term);
        term.c_cflag &= ~CLOCAL;
        ioctl(fd, TCSETA, &term);
        */
        alarm((4*strlen(argv[2])) + 5);
```

```
again:
        if( fgets(buffer,sizeof buffer,fdr) == (char *)NULL )
                        exit(-1);
        if( strncmp(buffer, "NO CARRIER",10) == SAME ) {
                exit(-1);
        }
        if( strncmp(buffer, "CONNECT",7) != SAME ) {
                goto again;
        }
        exit(0);
}

alrmint()
{
        exit(-1);
}
```

Appendix D
M4: A Macro Processor

D.1 Introduction D-1

D.2 Invoking *m4* D-1

D.3 Defining Macros D-2

D.4 Quoting D-3

D.5 Using Arguments D-5

D.6 Using Arithmetic Built-ins D-6

D.7 Manipulating Files D-7

D.8 Using System Commands D-7

D.9 Using Conditionals D-8

D.10 Manipulating Strings D-8

D.11 Printing D-10

D.1 Introduction

The **m4** macro processor defines and processes specially defined strings of characters called macros. By defining a set of macros to be processed by **m4**, a programming language can be enhanced to make it:

— More structured

— More readable

— More appropriate for a particular application

The **#define** statement in C and the analogous **define** in Ratfor are examples of the basic facility provided by any macro processor—replacement of text by other text.

Besides the straightforward replacement of one string of text by another, **m4** provides:

— Macros with arguments

— Conditional macro expansions

— Arithmetic expressions

— File manipulation facilities

— String processing functions

The basic operation of **m4** is copying its input to its output. As the input is read, each alphanumeric token (that is, string of letters and digits) is checked. If the token is the name of a macro, then the name of the macro is replaced by its defining text. The resulting string is reread by **m4**. Macros may also be called with arguments, in which case the arguments are collected and substituted in the right places in the defining text before **m4** rescans the text.

M4 provides a collection of about twenty built-in macros. In addition, the user can define new macros. Built-ins and user-defined macros work in exactly the same way, except that some of the built-in macros have side effects on the state of the process.

D.2 Invoking *m4*

The invocation syntax for **m4** is:

m4 [*files*]

Each filename argument is processed in order. If there are no arguments, or if

an argument is a dash (–), then the standard input is read. The processed text is written to the standard output, and can be redirected as in the following example:

```
m4 file1 file2 - >outputfile
```

Note the use of the dash in the above example to indicate processing of the standard input, *after* the files *file1* and *file2* have been processed by **m4**.

D.3 Defining Macros

The primary built-in function of **m4** is **define**, which is used to define new macros. The input

```
define(name, stuff)
```

causes the string *name* to be defined as *stuff*. All subsequent occurrences of *name* will be replaced by *stuff*. *Name* must be alphanumeric and must begin with a letter (the underscore (_) counts as a letter). *Stuff* is any text, including text that contains balanced parentheses; it may stretch over multiple lines.

Thus, as a typical example

```
define(N, 100)
.
.
.
if (i > N)
```

defines "N" to be **100**, and uses this symbolic constant in a later **if** statement.

The left parenthesis must immediately follow the word **define**, to signal that **define** has arguments. If a macro or built-in name is not followed immediately by a left parenthesis, "(", it is assumed to have no arguments. This is the situation for "N" above; it is actually a macro with no arguments. Thus, when it is used, no parentheses are needed following its name.

You should also notice that a macro name is only recognized as such if it appears surrounded by nonalphanumerics. For example, in

```
define(N, 100)
...
if (NNN > 100)
```

the variable "NNN" is absolutely unrelated to the defined macro "N", even though it contains three N's.

Things may be defined in terms of other things. For example

```
define(N, 100)
define(M, N)
```

defines both M and N to be 100.

What happens if "N" is redefined? Or, to say it another way, is "M" defined as "N" or as 100? In **m4**, the latter is true, "M" is 100, so even if "N" subsequently changes, "M" does not.

This behavior arises because **m4** expands macro names into their defining text as soon as it possibly can. Here, that means that when the string "N" is seen as the arguments of **define** are being collected, it is immediately replaced by 100; it's just as if you had said

```
define(M, 100)
```

in the first place.

If this isn't what you really want, there are two ways out of it. The first, which is specific to this situation, is to interchange the order of the definitions:

```
define(M, N)
define(N, 100)
```

Now "M" is defined to be the string "N", so when you ask for "M" later, you will always get the value of "N" at that time (because the "M" will be replaced by "N" which, in turn, will be replaced by 100).

D.4 Quoting

The more general solution is to delay the expansion of the arguments of **define** by quoting them. Any text surrounded by single quotation marks ` and ´ is not expanded immediately, but has the quotation marks stripped off. If you say

```
define(N, 100)
define(M, `N')
```

the quotation marks around the "N" are stripped off as the argument is being collected, but they have served their purpose, and "M" is defined as the string "N", not 100. The general rule is that **m4** always strips off one level of single quotation marks whenever it evaluates something. This is true even outside of macros. If you want the word "define" to appear in the output, you have to quote it in the input, as in

```
`define' = 1;
```

As another instance of the same thing, which is a bit more surprising, consider redefining "N":

```
define(N, 100)
...
define(N, 200)
```

Perhaps regrettably, the "N" in the second definition is evaluated as soon as it's seen; that is, it is replaced by 100, so it's as if you had written

```
define(100, 200)
```

This statement is ignored by **m4**, since you can only define things that look like names, but it obviously doesn't have the effect you wanted. To really redefine "N", you must delay the evaluation by quoting:

```
define(N, 100)
...
define('N', 200)
```

In **m4**, it is often wise to quote the first argument of a macro.

If the forward and backward quotation marks (` and ´) are not convenient for some reason, the quotation marks can be changed with the built-in **changequote**. For example:

```
changequote([, ])
```

makes the new quotation marks the left and right brackets. You can restore the original characters with just

```
changequote
```

There are two additional built-ins related to **define**. The built-in **undefine** removes the definition of some macro or built-in:

```
undefine('N')
```

removes the definition of "N". Built-ins can be removed with **undefine**, as in

```
undefine('define')
```

but once you remove one, you can never get it back.

The built-in **ifdef** provides a way to determine if a macro is currently defined. For instance, pretend that either the word "xenix" or "unix" is defined according to a particular implementation of a program. To perform operations according to which system you have you might say:

```
ifdef('xenix', 'define(system,1)' )
ifdef('unix', 'define(system,2)' )
```

Don't forget the quotation marks in the above example.

Ifdef actually permits three arguments: if the name is undefined, the value of **ifdef** is then the third argument, as in

```
ifdef('xenix', on XENIX, not on XENIX)
```

D.5 Using Arguments

So far we have discussed the simplest form of macro processing — replacing one string by another (fixed) string. User-defined macros may also have arguments, so different invocations can have different results. Within the replacement text for a macro (the second argument of its **define**) any occurrence of **$*n*** will be replaced by the *n*th argument when the macro is actually used. Thus, the macro *bump*, defined as

```
define(bump, $1 = $1 + 1)
```

generates code to increment its argument by 1:

```
bump(x)
```

is

```
x = x + 1
```

A macro can have as many arguments as you want, but only the first nine are accessible, through **$1** to **$9**. (The macro name itself is **$0**.) Arguments that are not supplied are replaced by null strings, so we can define a macro *cat* which simply concatenates its arguments, like this:

```
define(cat, $1$2$3$4$5$6$7$8$9)
```

Thus

```
cat(x, y, z)
```

is equivalent to

```
xyz
```

The arguments **$4** through **$9** are null, since no corresponding arguments were provided.

Leading unquoted blanks, tabs, or newlines that occur during argument collection are discarded. All other white space is retained. Thus:

```
define(a,   b   c)
```

defines "a" to be "b c".

Arguments are separated by commas, but parentheses are counted properly, so a comma protected by parentheses does not terminate an argument. That is, in

```
define(a, (b,c))
```

there are only two arguments; the second is literally "(b,c)". And of course a bare comma or parenthesis can be inserted by quoting it.

D.6 Using Arithmetic Built-ins

M4 provides two built-in functions for doing arithmetic on integers. The simplest is **incr**, which increments its numeric argument by 1. Thus, to handle the common programming situation where you want a variable to be defined as one more than N, write

```
define(N, 100)
define(N1, 'incr(N)')
```

Then "N1" is defined as one more than the current value of "N".

The more general mechanism for arithmetic is a built-in called **eval**, which is capable of arbitrary arithmetic on integers. It provides the following operators (in decreasing order of precedence):

```
unary + and -
** or ^   (exponentiation)
* / %  (modulus)
+ -
== != < <= > >=
!      (not)
& or &&(logical and)
| or ||    (logical or)
```

Parentheses may be used to group operations where needed. All the operands of an expression given to **eval** must ultimately be numeric. The numeric value of a true relation (like 1>0) is 1, and false is 0. The precision in **eval** is implementation dependent.

As a simple example, suppose we want "M" to be "2**N+1". Then

```
define(N, 3)
define(M, 'eval(2**N+1)')
```

As a matter of principle, it is advisable to quote the defining text for a macro unless it is very simple indeed (say just a number); it usually gives the result you want, and is a good habit to get into.

D.7 Manipulating Files

You can include a new file in the input at any time by the built-in function **include**:

include(*filename*)

inserts the contents of *filename* in place of the **include** command. The contents of the file is often a set of definitions. The value of **include** (that is, its replacement text) is the contents of the file; this can be captured in definitions, etc.

It is a fatal error if the file named in **include** cannot be accessed. To get some control over this situation, the alternate form **sinclude** can be used; **sinclude** (for "silent include") says nothing and continues if it can't access the file.

It is also possible to divert the output of **m4** to temporary files during processing, and output the collected material upon command. **M4** maintains nine of these diversions, numbered 1 through 9. If you say

divert(n)

all subsequent output is put onto the end of a temporary file referred to as "n". Diverting to this file is stopped by another **divert** command; in particular, **divert** or **divert(0)** resumes the normal output process.

Diverted text is normally output all at once at the end of processing, with the diversions output in numeric order. It is possible, however, to bring back diversions at any time, that is, to append them to the current diversion.

undivert

brings back all diversions in numeric order, and **undivert** with arguments brings back the selected diversions in the order given. The act of undiverting discards the diverted stuff, as does diverting into a diversion whose number is not between 0 and 9 inclusive.

The value of **undivert** is not the diverted stuff. Furthermore, the diverted material is not rescanned for macros.

The built-in **divnum** returns the number of the currently active diversion. This is zero during normal processing.

D.8 Using System Commands

You can run any program in the local operating system with the **syscmd** built-in. For example,

```
syscmd(date)
```

runs the **date** command. Normally, **syscmd** would be used to create a file for a subsequent **include**.

To facilitate making unique file names, the built-in **maketemp** is provided, with specifications identical to the system function **mktemp**: a string of "XXXXX" in the argument is replaced by the process id of the current process.

D.9 Using Conditionals

There is a built-in called **ifelse** which enables you to perform arbitrary conditional testing. In the simplest form,

```
ifelse(a, b, c, d)
```

compares the two strings *a* and *b*. If these are identical, **ifelse** returns the string *c*; otherwise it returns *d*. Thus, we might define a macro called **compare** which compares two strings and returns "yes" or "no" if they are the same or different.

```
define(compare, 'ifelse($1, $2, yes, no)')
```

Note the quotation marks, which prevent too-early evaluation of **ifelse**.

If the fourth argument is missing, it is treated as empty.

ifelse can actually have any number of arguments, and thus provides a limited form of multi-way decision capability. In the input

```
ifelse(a, b, c, d, e, f, g)
```

if the string *a* matches the string *b*, the result is *c*. Otherwise, if *d* is the same as *e*, the result is *f*. Otherwise the result is *g*. If the final argument is omitted, the result is null, so

```
ifelse(a, b, c)
```

is *c* if *a* matches *b*, and null otherwise.

D.10 Manipulating Strings

The built-in **len** returns the length of the string that makes up its argument. Thus

```
len(abcdef)
```

is 6, and

```
len((a,b))
```

is 5.

The built-in **substr** can be used to produce substrings of strings. For example

```
substr(s,i,n)
```

returns the substring of *s* that starts at position *i* (origin zero), and is *n* characters long. If *n* is omitted, the rest of the string is returned, so

```
substr('now is the time', 1)
```

is

```
ow is the time
```

If *i* or *n* are out of range, various sensible things happen.

The command

```
index(s1,s2)
```

returns the index (position) in *s1* where the string *s2* occurs, or –1 if it doesn't occur. As with **substr**, the origin for strings is 0.

The built-in **translit** performs character transliteration.

```
translit(s, f, t)
```

modifies *s* by replacing any character found in *f* by the corresponding character of *t*. That is

```
translit(s, aeiou, 12345)
```

replaces the vowels by the corresponding digits. If *t* is shorter than *f*, characters that don't have an entry in *t* are deleted; as a limiting case, if *t* is not present at all, characters from *f* are deleted from *s*. So

```
translit(s, aeiou)
```

deletes vowels from "s".

There is also a built-in called **dnl** which deletes all characters that follow it up to and including the next newline. It is useful mainly for throwing away empty lines that otherwise tend to clutter up **m4** output. For example, if you say

```
define(N, 100)
define(M, 200)
define(L, 300)
```

the newline at the end of each line is not part of the definition, so it is copied into the output, where it may not be wanted. If you add **dnl** to each of these lines, the newlines will disappear.

Another way to achieve this, is

```
divert(-1)
        define(...)
        ...
divert
```

D.11 Printing

The built-in **errprint** writes its arguments out on the standard error file. Thus, you can say

```
errprint('fatal error')
```

Dumpdef is a debugging aid that dumps the current definitions of defined terms. If there are no arguments, you get everything; otherwise you get the ones you name as arguments. Don't forget the quotation marks.

Index

%H% keyword 5-13
%I% keyword 5-13
%M% keyword 5-14
.cshrc file A-1
.login file A-1
.logout file A-2
.DEFAULT target 4-5
.IGNORE target 4-5
.PRECIOUS target 4-5
.SILENT target 4-5
.SUFFIXES target 4-11

a.out file
  creating 2-8
  default debugging file 6-2
  default output 2-2
  groups and physical segments 2-20
  optimizing 2-9
  removing stack probes 2-11
  setting the stack size 2-19
  stripping symbols 2-10
  default output file 7-1
adb command
  syntax 6-1
  – p option 6-3
  – w option 6-3
  addresses 6-4
  addresses and memory maps 6-26
  breakpoint (:br) command 6-15
  change memory map commands 6-25
  command prompt 6-3
  continue (:co) command 6-16
  create memory map commands 6-25
  data formats 6-9
  default input commands 6-29
  delete breakpoints (:dl) command 6-17
  display (=) command 6-10
  display backtrace ($c) command 6-17
  display breakpoint ($b) command 6-15
  display data (/) command 6-11
  display external variables ($e) command 6-18
  display memory map ($m) command 6-23
  display registers ($r) command 6-18
adb command *(continued)*
  display text (?) command 6-11
  expressions 6-5
  integers 6-5
  kill (:k) command 6-17
  locate commands 6-32
  maximum offset ($s) command 6-28
  memory maps 6-23
  operators 6-7
  output width ($w) command 6-28
  patching files and memory 6-32
  run (:R) command 6-14
  run (:r) command 6-14
  scripts 6-27
  single-step (:s) command 6-16
  stopping 6-4
  stopping a program 6-16
  symbols 6-5
  variables 6-6
  write commands 6-32
  writing to files 6-3
admin command
  – a option 5-23
  – d option 5-16
  – e option 5-24
  – f option 5-14
  – h option 5-25
  – i option 5-5
  – m option 5-18
  – t option 5-19
  – y option 5-27
  – z option 5-26
alias command A-7
ar command 1-2
Archives
  creating 1-2
  randomizing 1-2
ARGSUSED directive 3-11
as command
  syntax 7-1
  – l option 7-1
  – o option 7-1
  80286 instructions 7-18
  addressing modes 7-19
  assembly listing 7-1
  assignment statements 7-5
  based index operands
  based operands 7-21
  branch instructions 7-17
  byte instructions 7-16

as command *(continued)*
comments 7-2
constants 7-2
diagnostics 7-23
direct address operands 7-20
directives 7-8
expression statements 7-5
expressions 7-6
identifiers 7-2
immediate operands 7-20
indexed operands 7-21
indirect address operands 7-22
input/output instructions 7-18
instruction mnemonics 7-13
intersegment instructions 7-18
keyword statments 7-6
labels 7-4
null statements 7-5
output 7-1
register operands 7-19
segment directives 7-10
segments 7-3
statements 7-4
string instructions 7-18
types 7-6
whitespace 7-2
.ascii directive 7-11
.asciz directive 7-11
.blkb directive 7-12
.blkw directive 7-12
.bss directive 7-10
.byte directive 7-12
.comm directive 7-10
.data directive 7-10
.double directive 7-9
.end directive 7-12
.even directive 7-9
.float directive 7-9
.globl directive 7-9
.insrt directive 7-10
.list directive 7-11
.nlist directive 7-11
.text directive 7-10
.word directive 7-12
Assembler See as command
Assembly language source files
See Source file
Assignments, checked 3-8

Binary file See a.out file
Binary operations, checked 3-6
Branch number 5-2
break command A-16, A-19
Break statement
unreachable 3-4
breaksw command A-19

C compiler 2-1
C language source file See Source file
old syntax checked 3-9
C program check 3-1
C-shell See csh command
cc command
default output file (a.out) 2-2
memory models 2-5
naming the output file 2-4
syntax 2-1
– c option 2-7
– D option 2-13
– F option 2-19
– i option 2-6
– I option 2-14
– K option 2-11
– l option 2-8
– L option 2-11
– Ml option 2-7
– Mm option 2-6
– Ms option 2-6
– ND option 2-20
– NGD option 2-20
– NGT option 2-20
– NM option 2-19
– NT option 2-20
– o option 2-4
– O option 2-9
– p option 2-12
– s option 2-10
– S option 2-11
– W option 2-16
– w option 2-16
– x option 2-10
– X option 2-14
cdc command 5-17
Constant in conditional context 3-9
continue command A-16
core file, default 6-2
csh command
syntax A-1
– n option A-21
– v option A-20
– x option A-21
.cshrc file A-1

csh command *(continued)*
.login file A-1
.logout file A-2
:e modifier A-17
:h modifier A-17
:r modifier A-17
:t modifier A-18
:x modifier A-22
alias command A-7, A-9
aliases A-7
appending output (>>) A-8
argument symbol ($) A-6
argv variable A-12
arithmetic operations A-14
background job A-9
background symbol (&) A-9
break command A-16, A-19
breaksw command A-19
combining output (>&) A-8
command editing A-6
command output substitution (`) A-22
command prompt symbol (%) A-2
command quoting A-8
command quoting (") A-21
command repetition A-10
command separation (;) A-8
command termination status A-15
comment (#) A-12
comment symbol (#) A-18
continue command A-16
editing symbol (^) A-6
else-if statement A-16
environment setting A-10
escape character (\) A-8
expressions A-14
file enquiries A-15
filename extension extraction A-17
filename extraction A-17
foreach command A-16
goto statement A-19
history command A-6, A-10, A-10
history invocation symbol (!) A-6
history list A-4
history listing A-6
history variable A-2
home directory symbol (~) A-7
home variable A-3
if statement A-16
ignoreeof variable A-1, A-3
input from files A-10
input redirection A-20

csh command *(continued)*
input redirection ($<) A-14
interrupt catching A-20
INTERRUPT key A-9
interrupts A-20
logging out A-1, A-2
logout command A-1, A-10
looping A-16
mail variable A-1
modifiers A-17
noclobber variable A-4
onintr statement A-20
output redirection A-8
overwrite prevention A-4
overwrite symbol (>!) A-4
parsing variables A-22
path variable A-2
pathname extraction A-17, A-17
pipe A-8
printing previous commands A-6
process number ($$) A-14
prompt (%) A-2
QUIT key A-9
redirection A-8
rehash command A-3, A-10
repeat command A-10
repeating commands A-6
script arguments A-12
script debugging A-20
script invocation A-11
scripts A-11
secondary prompt (?) A-21
set command A-2
setenv command A-10
source command A-10
standard error A-8
standard output A-8
status variable A-15
string comparisons A-14
substitution by elements A-13
substitution, repeating ({ }) A-22
substitution, suspending (') A-8
substitution, suspending (\) A-8
switch statement A-18
time command A-11
time variable A-2
timing commands A-11
unalias command A-11
unset command A-3, A-11
variable assignment A-2
variable listing A-2
variable number symbol ($#) A-12

csh command *(continued)*
variable substitution symbol ($) A-12
variable test symbol ($?) A-12
while statement A-18
ctags command 1-3

d-file 5-3
Debugger See adb command
Degenerate unsigned comparison 3-8
delta command 5-8
- l option 5-29
- lp option 5-29
- m option 5-18
- n option 5-12
- p option 5-30
- r option 5-22
- y option 5-17

Delta
defined 5-2
Dependency line
commands 4-1
dependent filename 4-1
pseudo-target names 4-4
syntax 4-1
target filename 4-1
Dependency lines
library as a target 4-12
Disambiguating rules 9-16
Display
program contents 1-3
Dump
hexadecimal 1-3
strings 1-3
symbol list 1-3

Environment
make command 4-8
Error messages
as command 2-15, 7-23
C compiler 2-15
C preprocessor 2-15
cc command 2-15
ld command 2-15
SCCS commands 5-12
Executable file See a.out file
Expression
evaluation order 3-10
External declarations 3-2

far keyword 2-17
File administrator 5-4
foreach command A-16
Functions
return values 3-5
types checked 3-6
unused 3-2
see also alphabetical listing XENIX Reference Manual

g-file 5-3
get command 5-6
concurrent editing 5-21
- e option 5-7
- g option 5-26
- i option 5-28
- k option 5-26
- m option 5-30
- n option 5-30
- p option 5-11
- r option 5-9
- s option 5-28
- t option 5-11
- x option 5-28

goto statement A-19
Group 2-20

hd command 1-3
help command 5-12
history command A-6

Identification keywords 5-13
Implied assignments, checked 3-6
Impure-text programs 2-5, 2-6
Include directories 2-14
Infinite loops 3-4

l-file 5-3
Large model programs 2-5, 2-7
Level number 5-2
lex command
syntax 8-4
- l option 8-5

output file 8-5
%Start keyword 8-16
action 8-8
actions 8-3
ambiguous source rules 8-12

lex command *(continued)*
arbitrary character (.) operator 8-6
BEGIN action 8-16
beginning of line (^) operator 8-7
buffer overflow 8-13
character class ([ ]) operator 8-5
character classes 8-5
character range (-) operator 8-5
character translation table 8-22
current text, yytext 8-9
default action 8-8
definition ({ }) operator 8-8
ECHO action 8-9
echoing a match 8-9
end of line ($) operator 8-7
end-of-file value (0) 8-12
entering a start condition 8-16
environment change 8-15
escape (\) operator 8-4
escape characters 8-4
expression operator list 8-4
flags 8-15
grouping (( )) operator 8-7
ignoring input 8-8
input() function 8-11
length of current text, yyleng 8-9
lex.yy.c file 8-5
literal character ('') operator 8-4
lookahead function, yymore() 8-10
matching occurrences 8-13
matching preferences 8-12
not (^) operator 8-5
octal-coded characters 8-6
optional character (?) operator 8-6
or (|) operator 8-7
output() function 8-11
regular expressions 8-3
REJECT action 8-14
rejecting a match 8-15
repeat action (|) character 8-9
repetition ({ }) operator 8-8
repetition (*) operator 8-6
repetition (+) operator 8-6
rescanning function, yyless() 8-10
rules of source format 8-17
source definitions 8-17
source file format 8-2
source file separator (%) 8-3
source format 8-17
start condition 8-15
start condition indicator (< >)
8-7, 8-16
lex command *(continued)*
substitution strings 8-18
trailing context (/) operator 8-7
unput() function 8-11
wrap up function, yywrap() 8-12
yacc interface, yylex() function 8-18
yyleng variable 8-9
yyless() function 8-10
yylex() function, yacc interface 8-18
yymore() function 8-10
yytext array 8-9
yywrap() function 8-12
Libraries
creating 1-2
lex library 8-5
linking 2-8
lint command 3-12
make command 4-12
ordering relation 1-2
random access 1-2
yacc library 9-24
lint command
directives 3-11
libraries 3-12
syntax 3-1
– a option 3-8
– b option 3-4
– c option 3-7
– h option 3-9
– n option 3-12
– p option 3-12
– u option 3-3
– v option 3-3
– x option 3-2
LINTLIBRARY directive 3-12
Loader See ld command
logout command A-1
Loops 3-4
lorder command 1-2

Macros
defining 2-13
make command 4-5
preprocessing 1-2
make command
built-in macros 4-7
built-in rules 4-9
debugging 4-13
dependency line 4-1
environment variables 4-8
macros 4-5

make command *(continued)*
  makefile syntax 4-1
  options 4-4
  program maintenance 4-1
  suffix rules 4-10
  syntax 4-3
  troubleshooting 4-13
makefile
  debugging 4-13
  dependency line syntax 4-1
  macro syntax 4-5
Memory models, described 2-5
Middle model programs 2-5, 2-6
Modification requests 5-18
Modules 2-19

near keyword 2-17
nm command 1-3
Nonportable characters 3-7
NOSTRICT directive 3-11
Notational conventions 1-5
NOTREACHED directive 3-11

Object file
  creating 2-4, 2-7
  group name 2-20
  linking 2-4
  module name 2-19
  name convention 2-4, 7-1
  optimizing 2-9
  segment name 2-19
  stripping symbols 2-10
Old C syntax, checked 3-9
Operator precedence 3-9

p-file 5-3
Pointer alignment, checked 3-10
Preprocessed file 2-14
Preprocessing 2-13
Profiling 2-12
Program development 1-1
Programs
  assembly language 1-2
  assembly language listing 2-11
  C language 1-1
  checking 1-1, 3-1
  creating 1-1
  debugging 1-1
  examining 1-3

Programs *(continued)*
  lexical analyzers 1-2
  memory models 2-5
  optimizing 2-9
  parsers 1-2
  preprocessing 2-13
  profiling 2-12
  removing stack probes 2-11
  shell 1-3
  source code control 1-3
  source file maintainer 1-2, 4-1
  stack size 2-19
  word order 2-18
prs command
  – d option 5-20
  – r option 5-20
Pure-text programs02-05, 2-6

q-file 5-4

ranlib command 1-2
rehash command A-3
Release number 5-2
rmdel command 5-31
Routines See alphabetical listing
    XENIX Reference Manual

s-file 5-3
SCCS
  %H% keyword 5-13
  %I% keyword 5-13
  %M% keyword 5-14
  adding comments 5-17
  c flag 5-24
  changing comments 5-17
  changing descriptive text 5-19
  changing release numbers 5-9
  command arguments 5-4
  comments 5-17, 5-27
  comparing s-files 5-32
  concurrent editing 5-21
  creating a branch number 5-10
  creating an s-file 5-5
  d flag 5-16
  d-file 5-4
  delta list 5-29
  delta table 5-17
  description field 5-17
  directory use 5-1

SCCS *(continued)*
displaying a version 5-11
editing a file 5-7
error messages 5-12
f flag 5-24
file administrator 5-4
file checking 5-25
file checksum 5-26
file protection 5-23
file use 5-1
flags 5-15
g-file 5-3
i flag 5-14, , 5-15
identification keyword 5-13
identification string 5-2
j flag 5-22
l flag 5-24
l-file 5-3
locking versions 5-24
m flag 5-14
modification requests 5-18
p-file 5-3
printing versions 5-20
purpose 5-1
q-file 5-4
regenerating a g-file 5-26
removing a version 5-31
removing descriptive text 5-19
removing flags 5-16
restoring a damaged p-file 5-26
retrieving a branch version 5-10
retrieving a file 5-6
retrieving a specific verison 5-9
retrieving the most recent version 5-11
s-file 5-3
saving a copy of a new version 5-12
saving a new file 5-8
searching for strings 5-31
setting flags 5-15
setting floor and ceiling 5-24
standard input 5-27
suppressing normal output 5-28
user list 5-23
v flag 5-16
x-file 5-3
z-file 5-3
sccsdiff command 5-32
Segment 2-19, 7-3
Sequence number 5-2
set command A-2
SID
SID *(continued)*
branch number 5-2
level number 5-2
release number 5-2
SCCS identification string 5-2
sequence number 5-2
size command 1-3
Small model programs02-05, 2-6
Software development 1-1
Source Code Control System See SCCS
Source file maintainer See Make
assembly language 2-17
assembly language listing 2-11
C language 2-2
compiling 2-2
name convention 2-2, 2-17, 7-1
preprocessing 2-13
Stack 2-19
Stack probes 2-11
strings command 1-3
strip command 1-3
Structures, types checked 3-6
Subroutines See alphabetical listing XENIX Reference Manual
Suffix rules 4-10
sum command 1-3
switch statement A-19
System calls See alphabetical listing XENIX Reference Manual

tsort command 1-2
Type casts, checked 3-7
Types
checking 3-6

Unreachable statement 3-4
unset command A-3
Unused functions 3-2
Unused variables 3-2

VARARGS directive 3-12
Variables
enumeration types checked 3-6
local 3-3
set and used 3-3
static 3-4
unused 3-2

Variables *(continued)*
what command 5-31
while statement A-18
Word order 2-18

x-file 5-3
xstr command 1-3

y.tab.c file 9-24
y.tab.h file 9-29
yacc command
  – d option 9-29
  – v option 9-12, 9-18
  %left keyword 9-19
  %noassoc keyword 9-19
  %prec keyword 9-20
  %right keyword 9-19
  %start keyword 9-6
  %token keyword 9-5
  %type keyword 9-30
  accept simulation 9-28
  actions 9-6
  ambiguities 9-14
  associativity 9-19
  comments 9-4
  conflicts 9-14
  context flags 9-27
  current state 9-10
  debugging information 9-25
  desk calculator 9-30
  disambiguating rules 9-16, 9-21
  endmarker 9-6, 9-9
  error handling 9-3, 9-22
  error recovery, yyerrok 9-23
  error response 9-22
  error simulation 9-28
  error token 9-9, 9-22
  escape character (\ ) 9-4
  escape sequences 9-5
  expressions, parsing 9-19
  finite state machine 9-10
  global declarations 9-7
  global definitions 9-7
  grammar rules 9-1, 9-4
  grammar rules, implied 9-7
  header file 9-29
  left recursion 9-26
  lexical analyzer 9-2, 9-8
  literals 9-4
  lookahead discard, yyclearin 9-23
yacc command *(continued)*
  lookahead token 9-10
  lookahead value 9-25
  main( ) function 9-24
  names 9-4
  nonterminal symbol 9-2
  nonterminal symbols 9-4
  old features 9-40
  parser
    accept action 9-12
    defined 9-1
    explanation 9-10
    output listing 9-12
    reduce action 9-11
    shift action 9-10
  precedence 9-19
  program requirements 9-24
  pseudo-variables ($) 9-6
  pseudo-variables, left context 9-28
  reduce/reduce conflict 9-16
  repetition symbol (|) 9-5
  reserved words 9-27
  return value types 9-29
  return values 9-6
  right recursion 9-26
  shift, reduce, accept, error actions 9-10
  shift/reduce conflict 9-16
  specification file format 9-4
  specification style 9-25
  stack 9-10
  start symbol 9-5
  terminal symbol 9-2
  token declaration 9-5
  token name 9-9
  token number 9-8
  token number, default 9-9
  token types 9-29
  tokens 9-1, 9-4
  type casting (< >) 9-30
  union declaration 9-29
  value stack 9-12
  y.output file 9-12, 9-19
  y.tab.c file 9-24
  y.tab.h file 9-29
  yacc library 9-24
  yy name convention 9-8
  YYACCEPT macro 9-28
  yychar variable 9-25
  yyclearin statemnt 9-23
  yydebug variable 9-25
  yyerrok statement 9-23

yacc command *(continued)*
YYERROR macro 9-28
yyerror() function 9-24
yylval variable 9-8
yyparse() function 9-24
YYSTYPE variable 9-29

z-file 5-3

Contents

Programming Commands (CP)

intro	Introduces XENIX Software Development commands.
adb	Invokes a general-purpose debugger.
admin	Creates and administers SCCS files.
ar	Maintains archives and libraries.
as	Invokes the XENIX assembler.
cb	Beautifies C programs.
cc	Invokes the C compiler.
cdc	Changes the delta commentary of an SCCS delta.
comb	Combines SCCS deltas.
config	Configures a XENIX system.
cref	Makes a cross-reference listing.
csh	Invokes a shell command interpreter with C-like syntax.
ctags	Creates a tags file.
delta	Makes a delta (change) to an SCCS file.
get	Gets a version of an SCCS file.
gets	Gets a string from the standard input.
hdr	Displays selected parts of object files.
help	Asks for help about SCCS commands.
ld	Invokes the link editor.
lex	Generates programs for lexical analysis.
lint	Checks C language usage and syntax.
lorder	Finds ordering relation for an object library.
m4	Invokes a macro processor.
make	Maintains, updates, and regenerates groups of programs.
mkstr	Creates an error message file from C source.
nm	Prints name list.
prof	Displays profile data.
prs	Prints an SCCS file.
ranlib	Converts archives to random libraries.
ratfor	Converts Rational FORTRAN into standard FORTRAN.
regcmp	Compiles regular expressions.
rmdel	Removes a delta from an SCCS file.
sact	Prints current SCCS file editing activity.
sccsdiff	Compares two versions of an SCCS file.
size	Prints the size of an object file.
spline	Interpolates smooth curve.

stackuse	Determines stack requirements for C programs.
strings	Finds the printable strings in an object file.
strip	Removes symbols and relocation bits.
time	Times a command.
tsort	Sorts a file topologically.
unget	Undoes a previous get of an SCCS file.
uucp, uulog	Copies files from XENIX to XENIX.
uux	Executes command on remote XENIX.
val	Validates an SCCS file.
xref	Cross-references C programs.
xstr	Extracts strings from C programs.
yacc	Invokes a compiler-compiler.

Name

intro – Introduces XENIX Software Development commands.

Description

This section describes use of the individual commands available in the XENIX Software Development System. Each individual command is labeled with the letters CP to distinguish it from commands available in the XENIX Timesharing and Text Processing Systems. These letters are used for easy reference from other documentation. For example, the reference *cc*(CP) indicates a reference to a discussion of the cc command in this section, where the letter "C" stands for "command" and the letter "P" stands for "Programming".

Syntax

Unless otherwise noted, commands described in this section accept options and other arguments according to the following syntax:

name [*options*] [*cmdarg*]

where:

name The filename or pathname of an executable file

option A single letter representing a command option By convention, most options are preceded with a dash. Option letters can sometimes be grouped together as in – **abcd** or alternatively they are specified individually as in – **a** – **b** – **c** – **d** . The method of specifying options depends on the syntax of the individual command. In the latter method of specifying options, arguments can be given to the options. For example, the – **f** option for many commands often takes a following filename argument.

cmdarg A pathname or other command argument *not* beginning with a dash. It may also be a dash alone by itself indicating the standard input.

See Also

getopt(C), getopt(S)

Diagnostics

Upon termination, each command returns 2 bytes of status, one supplied by the system and giving the cause for termination, and (in the

case of "normal" termination) one supplied by the program (see *wait*(S) and *exit*(S)). The former byte is 0 for normal termination; the latter is customarily 0 for successful execution and nonzero to indicate troubles such as erroneous parameters, or bad or inaccessible data. It is called variously "exit code", "exit status", or "return code", and is described only where special conventions are involved.

Notes

Not all commands adhere to the above syntax.

Name

adb – Invokes a general-purpose debugger.

Syntax

adb [– **w**] [– p prompt] [objfil [corefile]]

Description

Adb is a general purpose debugging program. It may be used to examine files and to provide a controlled environment for the execution of XENIX programs.

Objfil is normally an executable program file, preferably containing a symbol table; if not then the symbolic features of *adb* cannot be used although the file can still be examined. The default for *objfil* is **a.out**. *Corefile* is assumed to be a core image file produced after executing *objfil*; the default for *corefile* is **core**.

Requests to *adb* are read from the standard input and responses are to the standard output. If the – w option is present then both *objfil* and *corefile* are created if necessary and opened for reading and writing so that files can be modified using *adb*. The QUIT and INTERRUPT keys cause *adb* to return to the next command. The – p option defines the prompt string. It may be any combination of characters. The default is an asterisk (*).

In general requests to *adb* are of the form:

[*address*] [, *count*] [*command*] [;]

If *address* is present then *dot* is set to *address*. Initially *dot* is set to 0. For most commands *count* specifies how many times the command will be executed. The default *count* is 1. *Address* is a special expression having the form:

[*segment*:] *offset*

where *segment* gives the address of a specific text or data segment, and *offset* gives an offset from the beginning of that segment. If *segment* is not given, the last segment value given in a command is used.

The interpretation of an address depends on the context it is used in. If a subprocess is being debugged then addresses are interpreted in the usual way in the address space of the subprocess. For further details of address mapping see *Addresses*.

Expressions

. The value of *dot*.

\+ The value of *dot* incremented by the current increment.

^ The value of *dot* decremented by the current increment.

" The last *address* typed.

integer An octal number if *integer* begins with a 0; a hexadecimal number if preceded by **#** or **0x**; otherwise a decimal number.

integer.fraction
A 32-bit floating point number.

'cccc' The ASCII value of up to 4 characters. \ may be used to escape a '.

< name
The value of *name*, which is either a variable name or a register name. *Adb* maintains a number of variables (see *Variables*) named by single letters or digits. If *name* is a register name then the value of the register is obtained from the system header in *corefile*. The register names are **ax bx cx dx di si bp fl ip cs ds ss es sp**. The name **fl** refers to the status flags.

symbol A *symbol* is a sequence of upper or lower case letters, underscores or digits, not starting with a digit. The value of the *symbol* is taken from the symbol table in *objfil*. An initial _ or ~ will be prepended to *symbol* if needed.

_ symbol
In C, the 'true name' of an external symbol begins with _. It may be necessary to use this name to disinguish it from internal or hidden variables of a program.

(*exp*) The value of the expression *exp*.

Monadic operators

***** *exp* The contents of the location addressed by *exp*.

\- *exp* Integer negation.

~ *exp* Bitwise complement.

Dyadic operators

Dyadic operators are left-associative and are less binding than monadic operators.

e1+ *e2* Integer addition.

e1- *e2* Integer subtraction.

*e1***e2* Integer multiplication.

e1%*e2* Integer division.

e1&*e2* Bitwise conjunction.

e1|*e2* Bitwise disjunction.

e1^*e2* Remainder after division of *e1* by *e2*.

e1#*e2* *E1* rounded up to the next multiple of *e2*.

Commands

Most commands consist of a verb followed by a modifier or list of modifiers. The following verbs are available. (The commands '?' and '/' may be followed by '*'; see *Addresses* for further details.)

?*f* Locations starting at *address* in *objfil* are printed according to the format *f*.

/*f* Locations starting at *address* in *corefile* are printed according to the format *f*.

=*f* The value of *address* itself is printed in the styles indicated by the format *f*. (For i format '?' is printed for the parts of the instruction that reference subsequent words.)

A *format* consists of one or more characters that specify a style of printing. Each format character may be preceded by a decimal integer that is a repeat count for the format character. While stepping through a format *dot* is incremented temporarily by the amount given for each format letter. If no format is given then the last format is used. The format letters available are as follows:

o 2	Prints 2 bytes in octal. All octal numbers output by *adb* are preceded by 0.
O 4	Prints 4 bytes in octal.
q 2	Prints in signed octal.
Q 4	Prints long signed octal.
d 2	Prints in decimal.
D 4	Prints long decimal.

x 2	Prints 2 bytes in hexadecimal.
X 4	Prints 4 bytes in hexadecimal.
u 2	Prints as an unsigned decimal number.
U 4	Prints long unsigned decimal.
f 4	Prints the 32 bit value as a floating point number.
F 8	Prints double floating point.
b 1	Prints the addressed byte in octal.
c 1	Prints the addressed character.
C 1	Prints the addressed character using the following escape convention. Character values 000 to 040 are printed as an at-sign (@) followed by the corresponding character in the octal range 0100 to 0140. The at-sign character itself is printed as @@.
s *n*	Prints the addressed characters until a zero character is reached.
S *n*	Prints a string using the at-sign (@) escape convention. Here *n* is the length of the string including its zero terminator.
Y 4	Prints 4 bytes in date format (see *ctime*(S)).
i n	Prints as machine instructions. *n* is the number of bytes occupied by the instruction. This style of printing causes variables 1 and 2 to be set to the offset parts of the source and destination respectively.
a 0	Prints the value of *dot* in symbolic form. Symbols are checked to ensure that they have an appropriate type as indicated below. / local or global data symbol ? local or global text symbol = local or global absolute symbol
A 0	Prints the value of *dot* in absolute form.
p 2	Prints the addressed value in symbolic form using the same rules for symbol lookup as a.
t 0	When preceded by an integer, tabs to the next appropriate tab stop. For example, 8t moves to the next 8-space tab stop.
r 0	Prints a space.
n 0	Prints a newline.
"..." 0	Prints the enclosed string.
^	Decrements *dot* by the current increment. Nothing is printed.
+	Increments *dot* by 1. Nothing is printed.
–	Decrements *dot* by 1. Nothing is printed.

newline
: If the previous command temporarily incremented *dot*, makes the increment permanent. Repeat the previous command with a *count* of 1.

[?/]l *value mask*
: Words starting at *dot* are masked with *mask* and compared with *value* until a match is found. If L is used then the match is for

4 bytes at a time instead of 2. If no match is found then *dot* is unchanged; otherwise *dot* is set to the matched location. If *mask* is omitted then – 1 is used.

[? /] w *value* ...
Writes the 2-byte *value* into the addressed location. If the command is W, writes 4 bytes. Odd addresses are not allowed when writing to the subprocess address space.

[? /] m *segnum fpos size*
Sets new values for the given segment's file position and size. If *size* is not given, then only the file position is changed. The *segnum* must the segment number of a segment already in the memory map (see *Addresses*). If ? is given, a text segment is affected; if / a data segment.

[? /] M *segnum fpos size*
Creates a new segment in the memory map. The segment is given file position *fpos* and physical size *size* . The *segnum* must not already exist in the memory map. If ? is given, a text segment is created; if / a data segment.

>*name*
Dot is assigned to the variable or register named.

! A shell is called to read the rest of the line following '!'.

$*modifier*
Miscellaneous commands. The available *modifiers* are:

- <*f* Read commands from the file *f* and return.
- >*f* Send output to the file *f*, which is created if it does not exist.
- r Print the general registers and the instruction addressed by ip. *Dot* is set to ip.
- f Print the floating registers in single or double length.
- b Print all breakpoints and their associated counts and commands.
- c C stack backtrace. If *address* is given then it is taken as the address of the current frame (instead of bp). If C is used then the names and (16 bit) values of all automatic and static variables are printed for each active function. If *count* is given then only the first *count* frames are printed.
- e The names and values of external variables are printed.
- w Set the page width for output to *address* (default 80).
- s Set the limit for symbol matches to *address* (default 255).
- o Sets input and output default format to octal.
- d Sets input and output default format to decimal.
- x Sets input and output default format to hexadecimal.
- q Exit from *adb*.
- v Print all non zero variables in octal.
- m Print the address map.

:*modifier*

Manage a subprocess. Available modifiers are:

brc Set breakpoint at *address*. The breakpoint is executed *count*– 1 times before causing a stop. Each time the breakpoint is encountered the command *c* is executed. If this command sets *dot* to zero then the breakpoint causes a stop.

dl Delete breakpoint at *address*.

r [*arguments*]
Run *objfil* as a subprocess. If *address* is given explicitly then the program is entered at this point; otherwise the program is entered at its standard entry point. *count* specifies how many breakpoints are to be ignored before stopping. *Arguments* to the subprocess may be supplied on the same line as the command. An argument starting with < or > causes the standard input or output to be established for the command. All signals are turned on on entry to the subprocess.

R [*arguments*]
Same as the **r** command except that *arguments* are passed through a shell before being passed to to the program. This means shell metacharacters can be used in filenames.

co*s* The subprocess is continued and signal *s* is passed to it, see *signal*(S). If *address* is given then the subprocess is continued at this address. If no signal is specified then the signal that caused the subprocess to stop is sent. Breakpoint skipping is the same as for **r**.

s*s* As for **co** except that the subprocess is single stepped *count* times. If there is no current subprocess then *objfil* is run as a subprocess as for **r**. In this case no signal can be sent; the remainder of the line is treated as arguments to the subprocess.

k The current subprocess, if any, is terminated.

Variables

Adb provides a number of variables. Named variables are set initially by *adb* but are not used subsequently. Numbered variables are reserved for communication as follows.

0 The last value printed.
1 The last offset part of an instruction source.
2 The previous value of variable 1.

On entry the following are set from the system header in the *corefile*. If *corefile* does not appear to be a core file then these values are set from *objfil*:

b The base address of the data segment.
d The data segment size.
e The entry point.
m The execution type.
n The number of segments.
s The stack segment size.
t The text segment size.

Addresses

Addresses in *adb* refer to either a location in a file or in actual memory. When there is no current process in memory, *adb* addresses are computed as file locations, and requested text and data are read from the *objfil* and *corefile* files. When there is a process, such as after a :r command, addresses are computed as actual memory locations.

All text and data segments in a program have associated memory map entries. Each entry has a unique *segment number*. In addition, each entry has the *file position* of that segment's first byte, and the *physical size* of the segment in the file. When a process is running, a segment's entry has a *virtual size* which defines the size of the segment in memory at the current time. This size can change during execution.

When a address is given and no process is running, the file location corresponding to the address is calculated as:

effective-file-address = file-position + offset

If a process is running, the memory location is simply the offset in the given segment. These addresses are valid if and only if

0 <= offset <= size

where *size* is physical size for file locations and virtual size for memory locations. Otherwise, the requested *address* is not legal.

The initial setting of both mappings is suitable for normal a.out and core files. If either file is not of the kind expected then, for that file, *file position* is set to 0, and *size* is set to the maximum file size. In this way, the whole file can be examined with no address translation.

So that *adb* may be used on large files all appropriate values are kept as signed 32 bit integers.

Files

/dev/mem
/dev/swap
a.out
core

See Also

ptrace(S), a.out(F), core(F)

Diagnostics

The message "adb" appears when there is no current command or format.

Comments about inaccessible files, syntax errors, abnormal termination of commands, etc.

Exit status is 0, unless last command failed or returned nonzero status.

Notes

A breakpoint set at the entry point is not effective on initial entry to the program.

System calls cannot be single stepped.

Local variables whose names are the same as an external variable may foul up the accessing of the external.

Name

admin – Creates and administers SCCS files.

Syntax

admin [– **n**] [– **i**[name]] [– **r**rel] [– **t**[name]] [– **f**flag[flag-val]] [– **d**flag[flag-val]] [– **a**login] [– **e**login] [– **m**[mrlist]] [– **y**[comment]] [– **h**] [– **z**] files

Description

Admin is used to create new SCCS files and to change parameters of existing ones. Arguments to *admin*may appear in any order. They consist of options, which begin with – , and named files (note that SCCS filenames must begin with the characters **s.**). If a named file doesn't exist, it is created, and its parameters are initialized according to the specified options. Parameters not initialized by a option are assigned a default value. If a named file does exist, parameters corresponding to specified options are changed, and other parameters are left as is.

If a directory is named, *admin* behaves as though each file in the directory were specified as a named file, except that nonSCCS files (last component of the pathname does not begin with **s.**) and unreadable files are silently ignored. If the dash – is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed. Again, nonSCCS files and unreadable files are silently ignored.

The options are as follows. Each is explained as though only one named file is to be processed since the effects of the arguments apply independently to each named file.

– **n** This option indicates that a new SCCS file is to be created.

– **i**[*name*] The *name* of a file from which the text for a new SCCS file is to be taken. The text constitutes the first delta of the file (see – **r** below for delta numbering scheme). If the **i** option is used, but the filename is omitted, the text is obtained by reading the standard input until an end-of-file is encountered. If this option is omitted, then the SCCS file is created empty. Only one SCCS file may be created by an *admin* command on which the **i** option is supplied. Using a single *admin* to create two or more SCCS files require that they be created empty (no – **i** option). Note that the – **i** option implies the – **n** option.

– r*rel* The *rel*ease into which the initial delta is inserted. This option may be used only if the – i option is also used. If the – r option is not used, the initial delta is inserted into release 1. The level of the initial delta is always 1 (by default initial deltas are named 1.1).

– t[*name*] The *name* of a file from which descriptive text for the SCCS file is to be taken. If the – t option is used and *admin* is creating a new SCCS file (the – n and/or – i options also used), the descriptive text filename must also be supplied. In the case of existing SCCS files: a – t option without a filename causes removal of descriptive text (if any) currently in the SCCS file, and a – t option with a filename causes text (if any) in the named file to replace the descriptive text (if any) currently in the SCCS file.

– f*flag* This option specifies a *flag*, and possibly a value for the *flag*, to be placed in the SCCS file. Several f options may be supplied on a single *admin* command line. The allowable *flags* and their values are:

b Allows use of the – b option on a *get*(CP) command to create branch deltas.

c*ceil* The highest release (i.e., "ceiling"), a number less than or equal to 9999, which may be retrieved by a *get*(CP) command for editing. The default value for an unspecified c flag is 9999.

f*floor* The lowest release (i.e., "floor"), a number greater than 0 but less than 9999, which may be retrieved by a *get*(CP) command for editing. The default value for an unspecified f flag is 1.

d*SID* The default delta number (SID) to be used by a *get*(CP) command.

i Causes the "No id keywords (ge6)" message issued by *get*(CP) or *delta*(CP) to be treated as a fatal error. In the absence of this flag, the message is only a warning. The message is issued if no SCCS identification keywords (see *get*(CP)) are found in the text retrieved or stored in the SCCS file.

j Allows concurrent *get*(CP) commands for editing on the same SID of an SCCS file. This allows multiple concurrent updates to the same version of the SCCS file.

l*list* A *list* of releases to which deltas can no longer be made (get – e against one of these "locked" releases fails). The *list* has the following syntax:

<list> ::= <range> | <list> , <range>
<range> ::= *RELEASE NUMBER* | **a**

The character **a** in the *list* is equivalent to specifying *all releases* for the named SCCS file.

n Causes *delta*(CP) to create a "null" delta in each of those releases (if any) being skipped when a delta is made in a *new* release (e.g., in making delta 5.1 after delta 2.7, releases 3 and 4 are skipped). These null deltas serve as "anchor points" so that branch deltas may later be created from them. The absence of this flag causes skipped releases to be nonexistent in the SCCS file preventing branch deltas from being created from them in the future.

q*text* User-definable text substituted for all occurrences of the keyword in SCCS file text retrieved by *get*(CP).

m*mod* *Mod*ule name of the SCCS file substituted for all occurrences of the admin.CP keyword in SCCS file text retrieved by *get*(CP). If the **m** flag is not specified, the value assigned is the name of the SCCS file with the leading **s.** removed.

t*type* *Type* of module in the SCCS file substituted for all occurrences of
keyword in SCCS file text retrieved by *get*(CP).

v[*pgm*] Causes *delta*(CP) to prompt for Modification Request (MR) numbers as the reason for creating a delta. The optional value specifies the name of an MR number validity checking program (see *delta*(CP)). (If this flag is set when creating an SCCS file, the **m** option must also be used even if its value is null).

– d[*flag*] Causes removal (deletion) of the specified *flag* from an SCCS file. The – **d** option may be specified only when processing existing SCCS files. Several – **d** options may be supplied on a single *admin* command. See the – **f** option for allowable *flag* names.

l*list* A *list* of releases to be "unlocked". See the – f option for a description of the l flag and the syntax of a *list*.

– a*login* A *login* name, or numerical XENIX group ID, to be added to the list of users which may make deltas (changes) to the SCCS file. A group ID is equivalent to specifying all *login* names common to that group ID. Several a options may be used on a single *admin* command line. As many *login*s, or numerical group IDs, as desired may be on the list simultaneously. If the list of users is empty, then anyone may add deltas.

– e*login* A *login* name, or numerical group ID, to be erased from the list of users allowed to make deltas (changes) to the SCCS file. Specifying a group ID is equivalent to specifying all *login* names common to that group ID. Several e options may be used on a single *admin* command line.

– y[*comment*] The *comment* text is inserted into the SCCS file as a comment for the initial delta in a manner identical to that of *delta*(CP). Omission of the – y option results in a default comment line being inserted in the form:

YY/MM/DD HH:MM:SS by *login*

The – y option is valid only if the – i and/or – n options are specified (i.e., a new SCCS file is being created).

– m[*mrlist*] The list of Modification Requests (MR) numbers is inserted into the SCCS file as the reason for creating the initial delta in a manner identical to *delta*(CP). The v flag must be set and the MR numbers are validated if the v flag has a value (the name of an MR number validation program). Diagnostics will occur if the v flag is not set or MR validation fails.

– h Causes *admin* to check the structure of the SCCS file (see *sccsfile*(F)), and to compare a newly computed checksum (the sum of all the characters in the SCCS file except those in the first line) with the checksum that is stored in the first line of the SCCS file. Appropriate error diagnostics are produced.

This option inhibits writing on the file, nullifying the effect of any other options supplied, and is therefore only meaningful when processing existing files.

– z The SCCS file checksum is recomputed and stored in the first line of the SCCS file (see **– h**, above).

Note that use of this option on a truly corrupted file may prevent future detection of the corruption.

Files

The last component of all SCCS filenames must be of the form **s.***file-name*. New SCCS files are created read-only (**444** modified by umask) (see *chmod*(C)). Write permission in the pertinent directory is, of course, required to create a file. All writing done by *admin* is to a temporary x-file, called **x.***filename*, (see *get*(CP)), created with read-only permission if the *admin* command is creating a new SCCS file, or with the same mode as the SCCS file if it exists. After successful execution of *admin*, the SCCS file is removed (if it exists), and the x-file is renamed with the name of the SCCS file. This ensures that changes are made to the SCCS file only if no errors occurred.

It is recommended that directories containing SCCS files be mode 755 and that SCCS files themselves be read-only. The mode of the directories allows only the owner to modify SCCS files contained in the directories. The mode of the SCCS files prevents any modification at all except by SCCS commands.

If it should be necessary to patch an SCCS file for any reason, the mode may be changed to 644 by the owner allowing use of a text editor. *Care must be taken!* The edited file should *always* be processed by an **admin – h** to check for corruption followed by an **admin – z** to generate a proper checksum. Another **admin – h** is recommended to ensure the SCCS file is valid.

Admin also makes use of a transient lock file (called **z.***filename*), which is used to prevent simultaneous updates to the SCCS file by different users. See *get*(CP) for further information.

See Also

delta(CP), ed(C), get(CP), help(CP), prs(CP), what(C), sccsfile(F)

Diagnostics

Use *help*(CP) for explanations.

Name

ar – Maintains archives and libraries.

Syntax

ar key [posname] afile name ...

Description

Ar maintains groups of files combined into a single archive file. Its main use is to create and update library files as used by the link editor though it can be used for any similar purpose.

Key is one character from the set **drqtpmx**, optionally concatenated with one or more of **vuaibcln**. *Afile* is the archive file. The *names* are constituent files in the archive file. The *posname* is the name of a constituent file, and is required when certain keys are used. The meanings of the *key* characters are:

d Deletes the named files from the archive file.

r Replaces the named files in the archive file. If the optional character **u** is used with **r**, then only those files with modified dates later than the archive files are replaced. If an optional positioning character from the set **abi** is used, then the *posname* argument must be present and specifies that new files are to be placed after (**a**) or before (**b** or **i**) *posname*. Otherwise new files are placed at the end.

q Quickly appends the named files to the end of the archive file. Optional positioning characters are invalid. The command does not check whether the added members are already in the archive. Useful only to avoid quadratic behavior when creating a large archive piece by piece.

t Prints a table of contents of the archive file. If no names are given, all files in the archive are tabled. If names are given, only those files are tabled.

p Prints the named files in the archive.

m Moves the named files to the end of the archive. If a positioning character is present, then the *posname* argument must be present and, as in **r**, specifies where the files are to be moved.

x Extracts the named files. If no names are given, all files in the archive are extracted. Unless the optional character **n** is used with **x**, an extracted file's modification date will be set to the date stored in that file's archive header. In neither case does **x**

alter the archive file.

v Verbose. Under the verbose option, *ar* gives a file-by-file description of the making of a new archive file from the old archive and the constituent files. When used with **t**, it gives a long listing of all information about the files. When used with **x**, it precedes each file with a name.

c Create. Normally *ar* will create *afile* when it needs to. The create option suppresses the normal message that is produced when *afile* is created.

l Local. Normally *ar* places its temporary files in the directory /tmp. This option causes them to be placed in the local directory.

n New. When used with the *key* character **x** it sets the extracted file's modification date to the current date.

When *ar* creates an archive, it always creates the header in the format of the local system (see *ar*(F)).

Files

/tmp/v* Temporary files

See Also

ld(CP), lorder(CP), ar(F)

Notes

If the same file is mentioned twice in an argument list, it may be put in the archive twice.

Name

as – Invokes the XENIX assembler.

Syntax

as [*options*] *file*.s

Description

As is the XENIX 8086/286 assembler. It reads and assembles 8086/286 assembly language instructions from the source file named *file*.s, and creates either a linkable object file named *file*.o, or an executable program named **a.out**. The extension .s is recommended but not required. If this extension is not given, *as* displays a warning and continues processing.

There are the following options:

– l Creates an assembly listing file named *file*.L. The file lists the source instructions, the assembled (binary) code for each instruction, and any assembly errors.

– nl *num*
Sets the maximum length of external symbols to *num* . Names longer than *num* are truncated before being copied to the external symbol table.

– g Directs the assembler to interpret undefined symbols as globally-defined external symbols. If not given, undefined symbols cause an assembly error.

– Mm
Creates a middle model object file suitable for linking with other middle model object files. The resulting text segment is name "*module*_TEXT" and module is named "*file*" where *file* is in uppercase letters.

– NT *name*
Sets the text segment name of the assembled code to *name* . This option overrides the default text segment.

– NM *name*
Sets the module name of the assembled code to *name* . The

option overrides the default module name.

– o *objfile*
Copies the assembled instructions to the file named *objfile*. This file is executable only if no errors occur during the assembly. This option overrides the default object file name.

Files

/bin/as

See Also

a.out(F), cc(CP), ld(CP)

Notes

This assembler can assemble all instructions that are common to both the 8086 and 80286 instruction sets. This assembler cannot assemble instructions specific to the 80286.

Name

cb – Beautifies C programs.

Syntax

cb [file]

Description

Cb places a copy of the C program in *file* (standard input if *file* is not given) on the standard output with spacing and indentation that displays the structure of the program.

Name

cc – Invokes the C compiler.

Syntax

cc [*options*] *filename* ...

Description

Cc is the XENIX C compiler command. It creates executable programs by compiling and linking the files named by the *filename* arguments. *Cc* copies the resulting program to the file **a.out** .

The *filename* can name any C or assembly language source file or any object or library file. C source files must have a ".c" filename extension. Assembly language source files must ".s", object files ".o", and library files ".a" extensions. *Cc* invokes the C compiler for each C source file and copies the result to an object file whose basename is the same as the source file but whose extension is ".o". *Cc* invokes the XENIX assembler, *as* , for each assembly source file and copies the result to an object file with extension ".o". *Cc* ignores object and library files until all source files have been compiled or assembled. It then invokes the XENIX link editor, *ld* , and combines all the object files it has created together with object files and libraries given in the command line to form a single program.

Files are processed in the order they are encountered in the command line, so the order of files is important. Library files are examined only if functions referenced in previous files have not yet been defined. Library files must be in *ranlib*(CP) format, that is, the first member must be named __.SYMDEF, which is a dictionary for the library. The library is searched repeatedly to satisfy as many references as possible. Only those functions that define unresolved references are concatenated. A number of "standard" libraries are searched automatically. These libraries support the standard C library functions and program startup routines. Which libraries are used depends on the program's memory model (see "Memory Models" below). The entry point of the resulting program is set to the beginning of the "main" program function.

There are the following options:

– **P**
Preprocesses each source file and copies the result to a file whose basename is the same as the source but whose extension is ".i". Preprocessing performs the actions specified by the preprocessing directives.

– **E**
Preprocesses each source file as described for – **P** , but copies

the result to the standard output. The option also places a #line directive with the current input line number and source file name at the beginning of output for each file.

– C
Preserves comments when preprocessing a file with – E or – P. That is, comments are not removed from the preprocessed source. This option may only be used in conjunction with – E or – P .

– D *name [= string]*
Defines *name* to the preprocessor as if defined by #define in each source file. The form "– D *name*" sets *name* to 1. The form "– D *name = string*" sets *name* to the given *string*.

– I *pathname*
Adds *pathname* to the list of directories to be searched when an #include file is not found in the directory containing the current source file or whenever angle brackets (< >) enclose the filename. If the file cannot be found in directories in this list, directories in a standard list are searched.

– X
Removes the standard directories from the list of directories to be searched for #include files.

– V *string*
Copies *string* to the object file created from the given source file. This option is often used for version control.

– W *num*
Sets the output level for compiler warning messages. If *num* is 0, no warning messages are issued. If 1, only warnings about program structure and overt type mismatches are issued. If 2, warnings about strong typing mismatches are issued. If 3, warnings for all automatic conversions are issued. This option does not affect compiler error message output.

– w
Prevents compiler warning messages from being issued. Same as "– W 0".

– p Adds code for program profiling. Profiling code counts the number of calls to each routine in the program and copies this information to the **mon.out** file. This file can be examined using the *prof*(CP) command.

– i Creates separate instruction and data spaces for small model programs. When the output file is executed, the program text and data areas are allocated separate physical segments. The text portion will be read-only and may be shared by all users executing the file. The option is implied when creating middle or large model program. (Not implemented on all machines.)

– F *num*
Sets the size of the program stack to *num* bytes. Default stack size if not given, is 2 Kbytes.

– K
Removes stack probes from a program. Stack probes are used to detect stack overflow on entry to program routines.

– nl *num*
Sets the maximum length of external symbols to *num*. Names longer than *num* are truncated before being copied to the external symbol table.

– M *string*
Sets the program configuration. This configuration defines the program's memory model, word order, data threshold. It also enables C language enhancements such as advanced instruction set and keywords. The *string* may be any combination of the following (the "s", "m", and "l" are mutually exclusive):

s	Creates a small model program (default).
m	Creates a middle model program.
l	Creates a large model program.
e	Enables the far and near keywords.
2	Enables 286 code generation for compiled C source files.
b	Reverses the word order for **long** types. High order word is first. Default is low order word first.
t *num*	Sets the size of the largest data item in the data group to *num*. Default is 32,767.

– c Creates a linkable object file for each source file but does not link these files. No executable program is created.

– o *filename*
Defines *filename* to be the name of the final executable program. This option overrides the default name **a.out**.

– l*library*
Searches *library* for unresolved references to functions. The *library* must be an object file archive library in **ranlib** format.

– O
Invokes the object code optimizer.

– S Creates an assembly source listing of the compiled C source file and copies this listing to the file whose basename is the same as the source but whose extension is ".s". This file is suitable for assembly using *as*(CP).

– L
Creates an assembler listing file containing assembled code and assembly source instructions. The listing is copied to the file whose basename is the same as the source but whose extension is ".L". This options suppresses the "– S" option.

– NM *name*
Sets the module name for each compiled or assembled source file to *name*. If not given, the filename of each source file is used.

– NT *name*
Sets the text segment name for each compiled or assembled source file to *name*. If not given, the name "*module*_TEXT" is used for middle model, and "_TEXT" for small model.

– ND *name*
Sets the data segment name for each compiled or assembled source file to *name*. If not given, the name "_DATA" is used.

– NGT *name*
Sets the text group name for each compiled or assembled source file to *name*. If not given, the name "IGROUP" is used.

– NGD *name*
Sets the data group name for each compiled or assembled source file to *name*. If not given, the name "DGROUP" is used.

Many options (or equivalent forms of these options) are passed to the link editor as the last phase of compilation. The "s", "m", and "l" configuration options are passed to specify memory requirements. The – i, – F, and – p are passed to specify other characteristics of the final program.

The – D and – I options may be used several times on the command line. The – D option must not define the same name twice. These options affect subsequent source files only.

Memory Models

Cc can create programs for three different memory models: small, middle, and large. In addition, small model programs can be pure or impure.

Impure-Text Small Model
These programs occupy one 64 Kbyte physical segment in which both text and data are combined. *Cc* creates impure small model programs by default. They can also be created using the "-Ms" option.

Pure-Text Small Model
These programs occupy two 64 Kbyte physical segments. Text and data are in separate segments. The text is read-only and may be shared by several processes at once. The maximum program size is 128 Kbytes. Pure small model programs are created using the "-i" and "-Ms" options.

Middle Model

These programs occupy several physical segments, but only one segment contains data. Text is divided among as many segments as required. Special call and returns are used to access functions in other segments. Text can be any size. Data must not exceed 64 Kbytes. Middle models programs are created using the "-Mm" option. These programs are always pure.

Large Model

These programs occupy several physical segments with both text and data in as many segments as required. Special calls and returns are used to access functions in other segments. Special addresses are used to access data in other segments. Text and data may be any size, but no data item may be larger than 64 Kbytes. Large model programs are created using the "-Ml" option. These programs are always pure.

Small, middle, and large model object files can only be linked with object and library files of the same model. It is not possible to combine small, medium, and large model object files in one executable program. *Cc* automatically selects the correct small, middle, or large versions of the standard libraries based on the configuration option. It is up to the user to make sure that all of his own object files and private libraries are properly compiled in the appropriate model.

The special calls and returns used in middle and large model programs may affect execution time. In particular, the execution time of a program which makes heavy use of functions and function pointers may differ noticably from small model programs.

In both middle and large model programs, function pointers are 32 bits long. In large model programs, data pointers are 32 bits long. Programs making use of such pointers must be written carefully to avoid incorrect declaration and use of these variables. *Lint*(CP) will help to check for correct use.

The – NM, – NT, – ND, – NGT, – NGD options may be used with middle and large model programs to direct the text and data of specific object files to named physical segments. All text having the same text segment name is placed in a single physical segment. Similarly, all data having the same data segment name is placed in a single physical segment.

Files

/bin/cc

See Also

as(CP), ar(CP), ld(CP), lint(CP), ranlib(CP)

Notes

Error messages are produced by the program that detects the error. These messages are usually produced by the C compiler, but may occasionally be produced by the assembler or the link loader.

All object module libraries must have a current *ranlib* directory.

Name

cdc – Changes the delta commentary of an SCCS delta.

Syntax

cdc – r*SID* [– **m**[*mrlist*]] [– **y**[*comment*]] *files*

Description

Cdc changes the delta commentary for the *SID* specified by the – **r** option, of each named SCCS file.

Delta commentary is defined to be the Modification Request (MR) and comment information normally specified via the *delta*(CP) command (– **m** and – **y** options).

If a directory is named, *cdc* behaves as though each file in the directory were specified as a named file, except that nonSCCS files (last component of the pathname does not begin with **s.**) and unreadable files are silently ignored. If a name of – is given, the standard input is read (see Warning); each line of the standard input is taken to be the name of an SCCS file to be processed.

Arguments to *cdc*, which may appear in any order, consist of options and file names.

All the described options apply independently to each named file:

– **r***SID* Used to specify the *S*CCS *ID*entification (*SID*) string of a delta for which the delta commentary is to be changed.

– **m**[*mrlist*] If the SCCS file has the **v** flag set (see *admin*(CP)) then a list of MR numbers to be added and/or deleted in the delta commentary of the *SID* specified by the – **r** option *may* be supplied. A null MR list has no effect.

MR entries are added to the list of MRs in the same manner as that of *delta*(CP). In order to delete an MR, precede the MR number with the character **!** (see Examples). If the MR to be deleted is currently in the list of MRs, it is removed and changed into a "comment" line. A list of all deleted MRs is placed in the comment section of the delta commentary and preceded by a comment line stating that they were deleted.

If – **m** is not used and the standard input is a terminal, the prompt **MRs?** is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. The **MRs?** prompt always precedes the **comments?** prompt (see – **y** option).

MRs in a list are separated by blanks and/or tab characters. An unescaped newline character terminates the MR list.

Note that if the **v** flag has a value (see *admin*(CP)), it is taken to be the name of a program (or shell procedure) which validates the correctness of the MR numbers. If a nonzero exit status is returned from the MR number validation program, *cdc* terminates and the delta commentary remains unchanged.

– **y**[*comment*] Arbitrary text used to replace the *comment*(s) already existing for the delta specified by the – **r** option. The previous comments are kept and preceded by a comment line stating that they were changed. A null *comment* has no effect.

If – **y** is not specified and the standard input is a terminal, the prompt "comments?" is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. An unescaped newline character terminates the *comment* text.

In general, if you made the delta, you can change its delta commentary; or if you own the file and directory you can modify the delta commentary.

Examples

The following:

```
cdc – r1.6 – m"bl78-12345 !bl77-54321 bl79-00001" – ytrouble
s.file
```

adds bl78-12345 and bl79-00001 to the MR list, removes bl77-54321 from the MR list, and adds the comment **trouble** to delta 1.6 of s.file.

The following interactive sequence does the same thing.

```
cdc – r1.6 s.file
MRs? !bl77-54321 bl78-12345 bl79-00001
comments? trouble
```

Warning

If SCCS file names are supplied to the *cdc* command via the standard input (– on the command line), then the **– m** and **– y** options must also be used.

Files

x-file See *delta*(CP)

z-file See *delta*(CP)

See Also

admin(CP), delta(CP), get(CP), help(CP), prs(CP), sccsfile(F)

Diagnostics

Use *help*(CP) for explanations.

Name

comb – Combines SCCS deltas.

Syntax

comb [– o] [– s] [– psid] [– clist] files

Description

Comb provides the means to combine one or more deltas in an SCCS file and make a single new delta. The new delta replaces the previous deltas, making the SCCS file smaller than the original.

Comb does not perform the combination itself. Instead, it generates a shell procedure that you must save and execute to reconstruct the given SCCS files. *Comb* copies the generated shell procedure to the standard output. To save the procedure, you must redirect the output to a file. The saved file can then be executed like any other shell procedure (see *sh*(C)).

When invoking *comb*, arguments may be specified in any order. All options apply to all named SCCS files. If a directory is named, *comb* behaves as though each file in the directory were specified as a named file, except that nonSCCS files (last component of the pathname does not begin with s.) and unreadable files are silently ignored. If a name of – is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed; nonSCCS files and unreadable files are silently ignored.

The options are as follows. Each is explained as though only one named file is to be processed, but the effects of any option apply independently to each named file.

– p*SID* The *SCCS ID*entification string (SID) of the oldest delta to be preserved. All older deltas are discarded in the reconstructed file.

– c*list* A *list* (see *get*(CP) for the syntax of a *list*) of deltas to be preserved. All other deltas are discarded.

– o For each **get** – e generated, this argument causes the reconstructed file to be accessed at the release of the delta to be created, otherwise the reconstructed file would be accessed at the most recent ancestor. Use of the – o option may decrease the size of the reconstructed SCCS file. It may also alter the shape of the delta tree of the original file.

– s This argument causes *comb* to generate a shell procedure that will produce a report for each file giving the filename, size (in blocks) after combining, original size (also in blocks), and percentage change computed by:

100 * (original – combined) / original

Before any SCCS files are actually combined, you should use this option to determine exactly how much space is saved by the combining process.

If no options are specified, *comb* will preserve only leaf deltas and the minimal number of ancestors needed to preserve the tree.

Files

comb????? Temporary files

See Also

admin(CP), delta(CP), get(CP), help(CP), prs(CP), sccsfile(F)

Diagnostics

Use *help*(CP) for explanations.

Notes

Comb may rearrange the shape of the tree of deltas. It may not save any space; in fact, it is possible for the reconstructed file to be larger than the original.

Name

config – configure a XENIX system

Syntax

/etc/config [– **t**] [– **l** file] [– **c** file] [– **m** file] dfile

Description

Config is a program that takes a description of a XENIX system and generates a file which is a C program defining the configuration tables for the various devices on the system.

The – **c** option specifies the name of the configuration table file; **c.c** is the default name.

The – **m** option specifies the name of the file that contains all the information regarding supported devices; **/etc/master** is the default name. This file is supplied with the XENIX system and should *not* be modified unless the user *fully* understands its construction.

The – **t** option requests a short table of major device numbers for character and block type devices. This can facilitate the creation of special files.

The user must supply *dfile*; it must contain device information for the user's system. This file is divided into two parts. The first part contains physical device specifications. The second part contains system-dependent information. Any line with an asterisk (*) in column 1 is a comment.

All configurations are assumed to have a set of required devices which must be present to run XENIX such as the system clock. These devices *must not* be specified in *dfile*.

First Part of *dfile*

Each line contains two fields, delimited by blanks and/or tabs in the following format:

devname number

where *devname* is the name of the device (as it appears in the **/etc/master** device table), and *number* is the number (decimal) of devices associated with the corresponding controller; *number* is optional, and if omitted, a default value which is the maximum value for that controller is used.

There are certain drivers that may be provided with the system, that are actually *pseudo-device* drivers; that is, there is no real hardware associated with the driver. Drivers of this type are identified on their respective manual entries.

Second Part of *dfile*

The second part contains three different types of lines. Note that *all* specifications of this part *are required*, although their order is arbitrary.

1. *Root/pipe device specification*

   Each line has three fields:

root	devname	minor
pipe	devname	minor

   where *minor* is the minor device number (in octal).

2. *Swap device specification*

   One line that contains five fields as follows:

swap	devname	minor	swplo	nswap

   where *swplo* is the lowest disk block (decimal) in the swap area and *nswap* is the number of disk blocks (decimal) in the swap area.

3. *Parameter specification*

   A number of lines of two fields each as follows (*number* is decimal):

buffers	number
inodes	number
files	number
mounts	number
swapmap	number
pages	number
calls	number
procs	number
maxproc	number
texts	number
clists	number
locks	number
timezone	number
daylight	0 or 1

Example

Suppose we wish to configure a system with the following devices:
one HD disk drive controller with 1 drive
one FD floppy disk drive controller with 1 driver

We must also specify the following parameter information:

root device is an HD (pseudo disk 3)
pipe device is an HD (pseudo disk 3)
swap device is an HD (pseudo disk 2)
with a swplo of 1 and an nswap of 2300
number of buffers is 50
number of processes is 50
maximum number of processes per user ID is 15
number of mounts is 8
number of inodes is 120
number of files is 120
number of calls is 30
number of texts is 35
number of character buffers is 150
number of swapmap entries is 50
number of memory pages is 512
number of file locks is 100
timezone is pacific time
daylight time is in effect

The actual system configuration would be specified as follows:

```
hd   1
fd   1
root   hd   3
pipe   hd   3
swap   hd   2     0   2300
* Comments may be inserted in this manner
buffers 50
procs          150
maxproc 15
mounts         8
inodes         120
files          120
calls          30
texts          35
clists         150
swapmap        50
pages (1024/2);
locks 100
timezone (8*60)
daylight 1
```

Files

/etc/master	default input master device table
c.c	default output configuration table file

See Also

master(F)

Diagnostics

Diagnostics are routed to the standard output and are self-explanatory.

Notes

The – t option does not know about devices that have aliases. However, the major device numbers are always correct.

Name

cref – Makes a cross-reference listing.

Syntax

cref [– **acilnostux123**] files

Description

Cref makes a cross-reference listing of assembler or C programs. The program searches the given *files* for symbols in the appropriate C or assembly language syntax.

The output report is in four columns:

1. Symbol
2. Filename
3. Current symbol or line number
4. Text as it appears in the file

Cref uses either an *ignore* file or an *only* file. If the – **i** option is given, the next argument is taken to be an *ignore* file; if the – **o** option is given, the next argument is taken to be an *only* file. *Ignore* and *only* files are lists of symbols separated by newlines. All symbols in an *ignore* file are ignored in columns 1 and 3 of the output. If an *only* file is given, only symbols in that file will appear in column 1. Only one of these options may be given; the default setting is – **i** using the default ignore file (see *FILES* below). Assembler predefined symbols or C keywords are ignored.

The – **s** option causes current symbols to be put in column 3. In the assembler, the current symbol is the most recent name symbol; in C, the current function name. The – **l** option causes the line number within the file to be put in column 3.

The – **t** option causes the next available argument to be used as the name of the intermediate file (instead of the temporary file **/tmp/crt??**). This file is created and is *not* removed at the end of the process.

The *cref* options are:

a Uses assembler format (default)

c Uses C format

i Uses an *ignore* file (see above)

l Puts line number in column 3 (instead of current symbol)

n Omits column 4 (no context)

o Uses an *only* file (see above)

s Current symbol in column 3 (default)

t User-supplied temporary file

u Prints only symbols that occur exactly once

x Prints only C external symbols

1 Sorts output on column 1 (default)

2 Sorts output on column 2

3 Sorts output on column 3

Files

/usr/lib/cref/* Assembler specific files

See Also

as(CP), cc(CP), sort(C), xref(CP)

Notes

Cref inserts an ASCII DEL character into the intermediate file after the eighth character of each name that is eight or more characters long in the source file.

Name

csh – Invokes a shell command interpreter with C-like syntax.

Syntax

csh [– **cefinstvVxX**] [arg ...]

Description

Csh is a command language interpreter. It begins by executing commands from the file .cshrc in the home directory of the invoker. If this is a login shell, then it also executes commands from the file .login there. In the normal case, the shell will then begin reading commands from the terminal, prompting with % . Processing of arguments and the use of the shell to process files containing command scripts will be described later.

The shell then repeatedly performs the following actions: a line of command input is read and broken into *words*. This sequence of words is placed on the command history list and then parsed. Finally each command in the current line is executed.

When a login shell terminates, it executes commands from the file .logout in the user's home directory.

Lexical structure

The shell splits input lines into words at blanks and tabs with the following exceptions. The characters &, | ;, <, >, (,), form separate words. If doubled in &&, || <<, or >>, these pairs form single words. These parser metacharacters may be made part of other words, or prevented their special meaning, by preceding them with \. A newline preceded by a \ is equivalent to a blank.

In addition strings enclosed in matched pairs of quotations, ', ` or ", form parts of a word; metacharacters in these strings, including blanks and tabs, do not form separate words. These quotations have semantics to be described subsequently. Within pairs of \ or " characters a newline preceded by a \ gives a true newline character.

When the shell's input is not a terminal, the character # introduces a comment which continues to the end of the input line. It does not have this special meaning when preceded by \ and placed inside the quotation marks `, ', and ".

Commands

A simple command is a sequence of words, the first of which specifies the command to be executed. A simple command or a sequence

of simple commands separated by | characters forms a pipeline. The output of each command in a pipeline is connected to the input of the next. Sequences of pipelines may be separated by ;, and are then executed sequentially. A sequence of pipelines may be executed without waiting for it to terminate by following it with an &. Such a sequence is automatically prevented from being terminated by a hangup signal; the *nohup* command need not be used.

Any of the above may be placed in parentheses to form a simple command (which may be a component of a pipeline, etc.) It is also possible to separate pipelines with || or && indicating, as in the C language, that the second is to be executed only if the first fails or succeeds respectively. (See *Expressions.)*

Substitutions

The following sections describe the various transformations the shell performs on the input in the order in which they occur.

History Substitutions

History substitutions can be used to reintroduce sequences of words from previous commands, possibly performing modifications on these words. Thus history substitutions provide a generalization of a *redo* function.

History substitutions begin with the character ! and may begin anywhere in the input stream if a history substitution is not already in progress. This ! may be preceded by a \ to prevent its special meaning; a ! is passed unchanged when it is followed by a blank, tab, newline, =, or (. History substitutions also occur when an input line begins with ^. This special abbreviation will be described later.

Any input line which contains history substitution is echoed on the terminal before it is executed as it could have been typed without history substitution.

Commands input from the terminal which consist of one or more words are saved on the history list, the size of which is controlled by the *history* variable. The previous command is always retained. Commands are numbered sequentially from 1.

For example, consider the following output from the history command:

```
 9  write michael
10  ex write.c
11  cat oldwrite.c
12  diff *write.c
```

The commands are shown with their event numbers. It is not usually necessary to use event numbers, but the current event number can be made part of the prompt by placing a ! in the prompt string.

With the current event 13 we can refer to previous events by event number !11, relatively as in !- 2 (referring to the same event), by a prefix of a command word as in !d for event 12 or !w for event 9, or by a string contained in a word in the command as in !?mic? also referring to event 9. These forms, without further modification, simply reintroduce the words of the specified events, each separated by a single blank. As a special case !! refers to the previous command; thus !! alone is essentially a *redo*. The form !# references the current command (the one being typed in). It allows a word to be selected from further left in the line, to avoid retyping a long name, as in !#:1.

To select words from an event we can follow the event specification by a : and a designator for the desired words. The words of a input line are numbered from 0, the first (usually command) word being 0, the second word (first argument) being 1, and so on. The basic word designators are:

0 First (command) word

n
*n*th argument

^ First argument, i.e. 1

$ Last argument

% Word matched by (immediately preceding) ?*s*? search

x- *y*
Range of words

- *y* Abbreviates 0- *y*

* Abbreviates ^- $, or nothing if only 1 word in event

*x** Abbreviates *x*- $

x-
Like *x** but omitting word $

The : separating the event specification from the word designator can be omitted if the argument selector begins with a ↑, $, *, - or %. After the optional word designator can be placed a sequence of modifiers, each preceded by a :. The following modifiers are defined:

h Removes a trailing pathname component

r Removes a trailing .xxx component

s/*l*/*r*/
Substitutes *l* for *r*

t Removes all leading pathname components

& Repeats the previous substitution

g Applies the change globally, prefixing the above

p Prints the new command but do not execute it

q Quotes the substituted words, preventing substitutions

x Like q, but breaks into words at blanks, tabs, and newlines

Unless preceded by a g the modification is applied only to the first modifiable word. In any case it is an error for no word to be applicable.

The left side of substitutions are not regular expressions in the sense of the editors, but rather strings. Any character may be used as the delimiter in place of /; a \ quotes the delimiter into the *l* and *r* strings. The character & in the right side is replaced by the text from the left. A \ quotes & also. A null *l* uses the previous string either from a *l* or from a contextual scan string *s* in !?*s*?. The trailing delimiter in the substitution may be omitted if a newline follows immediately as may the trailing ? in a contextual scan.

A history reference may be given without an event specification, e.g. !$. In this case the reference is to the previous command unless a previous history reference occurred on the same line in which case this form repeats the previous reference. Thus !?foo?^!$ gives the first and last arguments from the command matching ?foo?.

A special abbreviation of a history reference occurs when the first nonblank character of an input line is a ^. This is equivalent to !:s^, providing a convenient shorthand for substitutions on the text of the previous line. Thus ^lb^lib fixes the spelling of lib in the previous command. Finally, a history substitution may be surrounded with { and } if necessary to insulate it from the characters that follow. Thus, after ls – ld ~paul we might do !{l}a to do ls – ld ~paula, while !la would look for a command starting la.

Quotations With ´ and "

The quotation of strings by ´ and " can be used to prevent all or some of the remaining substitutions. Strings enclosed in ´ are prevented any further interpretation. Strings enclosed in " are variable and command expansion may occur.

In both cases, the resulting text becomes (all or part of) a single word; only in one special case (see ***Command Substitution*** below) does a " quoted string yield parts of more than one word; ´ quoted strings never do.

Alias Substitution

The shell maintains a list of aliases which can be established, displayed and modified by the *alias* and *unalias* commands. After a command line is scanned, it is parsed into distinct commands and the first word of each command, left-to-right, is checked to see if it has an alias. If it does, then the text which is the alias for that command is reread with the history mechanism available as though that command were the previous input line. The resulting words replace the command and argument list. If no reference is made to the history list, then the argument list is left unchanged.

Thus if the alias for ls is ls – l the command "ls /usr" would map to "ls – l /usr". Similarly if the alias for lookup was "grep !^ /etc/passwd" then "lookup bill" would map to "grep bill /etc/passwd".

If an alias is found, the word transformation of the input text is performed and the aliasing process begins again on the reformed input line. Looping is prevented if the first word of the new text is the same as the old by flagging it to prevent further aliasing. Other loops are detected and cause an error.

Note that the mechanism allows aliases to introduce parser metasyntax. Thus we can alias print "'pr \!* | lpr'" to make a command that paginates its arguments to the lineprinter.

Variable Substitution

The shell maintains a set of variables, each of which has as value a list of zero or more words. Some of these variables are set by the shell or referred to by it. For instance, the *argv* variable is an image of the shell's argument list, and words of this variable's value are referred to in special ways.

The values of variables may be displayed and changed by using the *set* and *unset* commands. Of the variables referred to by the shell a number are toggles; the shell does not care what their value is, only whether they are set or not. For instance, the *verbose* variable is a toggle which causes command input to be echoed. The setting of this variable results from the – v command line option.

Other operations treat variables numerically. The at-sign (@) command permits numeric calculations to be performed and the result assigned to a variable. However, variable values are always represented as (zero or more) strings. For the purposes of numeric operations, the null string is considered to be zero, and the second and subsequent words of multiword values are ignored.

After the input line is aliased and parsed, and before each command is executed, variable substitution is performed, keyed by dollar sign ($) characters. This expansion can be prevented by preceding the dollar sign with a backslash (\) except within double quotation marks

(") where it *always* occurs, and within single quotation marks (') where it *never* occurs. Strings quoted by back quotation marks (`) are interpreted later (see *Command substitution* below) so dollar sign substitution does not occur there until later, if at all. A dollar sign is passed unchanged if followed by a blank, tab, or end-of-line.

Input and output redirections are recognized before variable expansion, and are variable expanded separately. Otherwise, the command name and entire argument list are expanded together. It is thus possible for the first (command) word to generate more than one word, the first of which becomes the command name, and the rest of which become arguments.

Unless enclosed in double quotation marks or given the :q modifier, the results of variable substitution may eventually be command and filename substituted. Within double quotation marks (") a variable whose value consists of multiple words expands to a portion of a single word, with the words of the variable's value separated by blanks. When the :q modifier is applied to a substitution the variable expands to multiple words with each word separated by a blank and quoted to prevent later command or filename substitution.

The following sequences are provided for introducing variable values into the shell input. Except as noted, it is an error to reference a variable which is not set.

$name
${name}
: Are replaced by the words of the value of variable *name*, each separated by a blank. Braces insulate *name* from following characters which would otherwise be part of it. Shell variables have names consisting of up to 20 letters, digits, and underscores.

If *name* is not a shell variable, but is set in the environment, then that value is returned (but : modifiers and the other forms given below are not available in this case).

$name[selector]
${name[selector]}
: May be used to select only some of the words from the value of *name*. The selector is subjected to $ substitution and may consist of a single number or two numbers separated by a – . The first word of a variables value is numbered 1. If the first number of a range is omitted it defaults to 1. If the last member of a range is omitted it defaults to $#name. The selector * selects all words. It is not an error for a range to be empty if the second argument is omitted or in range.

$#name
${#name}
Gives the number of words in the variable. This is useful for later use in a [selector].

$0
Substitutes the name of the file from which command input is being read. An error occurs if the name is not known.

$number
${number}
Equivalent to $argv[number].

$* Equivalent to $argv[*].

The modifiers :h, :t, :r, :q and :x may be applied to the substitutions above as may :gh, :gt and :gr. If braces { } appear in the command form then the modifiers must appear within the braces. Only one : modifier is allowed on each $ expansion.

The following substitutions may not be modified with : modifiers.

$?name
${?name}
Substitutes the string 1 if name is set, 0 if it is not.

$?0 Substitutes 1 if the current input filename is known, 0 if it is not.

$$ Substitutes the (decimal) process number of the (parent) shell.

Command and Filename Substitution

Command and filename substitution are applied selectively to the arguments of built-in commands. This means that portions of expressions which are not evaluated are not subjected to these expansions. For commands which are not internal to the shell, the command name is substituted separately from the argument list. This occurs very late, after input-output redirection is performed, and in a child of the main shell.

Command Substitution

Command substitution is indicated by a command enclosed in back quotation marks. The output from such a command is normally broken into separate words at blanks, tabs and newlines, with null words being discarded, this text then replacing the original string. Within double quotation marks, only newlines force new words; blanks and tabs are preserved.

In any case, the single final newline does not force a new word. Note that it is thus possible for a command substitution to yield only part of a word, even if the command outputs a complete line.

Filename Substitution

If a word contains any of the characters *, ?, [or {, or begins with the character ~, then that word is a candidate for filename substitution, also known as globbing. This word is then regarded as a pattern, and replaced with an alphabetically sorted list of filenames which match the pattern. In a list of words specifying filename substitution it is an error for no pattern to match an existing filename, but it is not required for each pattern to match. Only the metacharacters *, ?, and [imply pattern matching, the characters ~ and { being more akin to abbreviations.

In matching filenames, the character . at the beginning of a filename or immediately following a /, as well as the character / must be matched explicitly. The character * matches any string of characters, including the null string. The character ? matches any single character. The sequence [...] matches any one of the characters enclosed. Within [...], a pair of characters separated by – matches any character lexically between the two.

The character ~ at the beginning of a filename is used to refer to home directories. Standing alone it expands to the invoker's home directory as reflected in the value of the variable *home*. When followed by a name consisting of letters, digits and – characters the shell searches for a user with that name and substitutes their home directory; thus ~ken might expand to /usr/ken and ~ken/chmach to /usr/ken/chmach. If the character ~ is followed by a character other than a letter or /, or appears not at the beginning of a word, it is left unchanged.

The metanotation a{b,c,d}e is a shorthand for abe ace ade. Left to right order is preserved, with results of matches being sorted separately at a low level to preserve this order. This construct may be nested. Thus ~source/s1/{oldls,ls}.c expands to /usr/source/s1/oldls.c /usr/source/s1/ls.c, whether or not these files exist, without any chance of error if the home directory for source is /usr/source. Similarly ../{memo,*box} might expand to ../memo ../box ../mbox. (Note that memo was not sorted with the results of matching *box.) As a special case {, } and {} are passed unchanged.

Input/Output

The standard input and standard output of a command may be redirected with the following syntax:

< name
: Opens file *name* (which is first variable, command and filename expanded) as the standard input.

<< word
: Reads the shell input up to a line which is identical to *word*. *Word* is not subjected to variable, filename or command substitution, and each input line is compared to *word* before any

substitutions are done on this input line. Unless a quoting backslash, double, or single quotation mark, or a back quotation mark appears in *word*, variable and command substitution is performed on the intervening lines, allowing \ to quote **$**, \ and `. Commands which are substituted have all blanks, tabs, and newlines preserved, except for the final newline which is dropped. The resulting text is placed in an anonymous temporary file which is given to the command as standard input.

> name
>! name
>& name
>&! name

The file *name* is used as standard output. If the file does not exist then it is created; if the file exists, it is truncated, and its previous contents is lost.

If the variable *noclobber* is set, then the file must not already exist or it must be a character special file (e.g. a terminal or /dev/null) or an error results. This helps prevent accidental destruction of files. In this case, the ! forms can be used to suppress this check.

The forms involving & route the diagnostic output into the specified file as well as the standard output. *Name* is expanded in the same way as < input filenames are.

>> name
>>& name
>>! name
>>&! name

Uses file *name* as standard output like > but places output at the end of the file. If the variable *noclobber* is set, then it is an error for the file not to exist unless one of the ! forms is given. Otherwise similar to >.

If a command is run detached (followed by &) then the default standard input for the command is the empty file /dev/null. Otherwise the command receives the environment in which the shell was invoked as modified by the input-output parameters and the presence of the command in a pipeline. Thus, unlike some previous shells, commands run from a file of shell commands have no access to the text of the commands by default; rather they receive the original standard input of the shell. The << mechanism should be used to present inline data. This permits shell command scripts to function as components of pipelines and allows the shell to block read its input.

Diagnostic output may be directed through a pipe with the standard output. Simply use the form |& rather than just |

Expressions

A number of the built-in commands (to be described later) take expressions, in which the operators are similar to those of C, with the same precedence. These expressions appear in the **@**, ***exit***, ***if***, and ***while*** commands. The following operators are available:

```
|| && | ^ & == != <= >= < > << >>
+ - * / % ! ~ ( )
```

Here the precedence increases to the right, with the operators:

```
== and !=
<=, >=, <, and >
<< and >>
+ and -
* / and %
```

forming groups at the same level. The == and != operators compare their arguments as strings, all others operate on numbers. Strings which begin with **0** are considered octal numbers. Null or missing arguments are considered **0**. The result of all expressions are strings, which represent decimal numbers. It is important to note that no two components of an expression can appear in the same word; except when adjacent to components of expressions which are syntactically significant to the parser (& | < > ( )) they should be surrounded by spaces.

Also available in expressions as primitive operands are command executions enclosed in { and } and file enquiries of the form **-** *l* name where *l* is one of:

r	Read access
w	Write access
x	Execute access
e	Existence
o	Ownership
z	Zero size
f	Plain file
d	Directory

The specified name is command and filename expanded, then tested to see if it has the specified relationship to the real user. If the file does not exist or is inaccessible then all enquiries return false, i.e. **0**. Command executions succeed, returning true, i.e. **1**, if the command exits with status **0**, otherwise they fail, returning false, i.e. **0**. If more detailed status information is required then the command should be executed outside of an expression and the variable ***status*** examined.

Control Flow

The shell contains a number of commands which can be used to regulate the flow of control in command files (shell scripts) and (in limited but useful ways) from terminal input. These commands all operate by forcing the shell to reread or skip in its input and, due to the implementation, restrict the placement of some of the commands.

The *foreach, switch,* and *while* statements, as well as the *if– then– else* form of the *if* statement require that the major keywords appear in a single simple command on an input line as shown below.

If the shell's input is not seekable, the shell buffers up input whenever a loop is being read and performs seeks in this internal buffer to accomplish the rereading implied by the loop. (To the extent that this allows, backward goto commands will succeed on nonseekable inputs.)

Built-In Commands

Built-in commands are executed within the shell. If a built-in command occurs as any component of a pipeline except the last then it is executed in a subshell.

alias
alias name
alias name wordlist
The first form prints all aliases. The second form prints the alias for *name* . The final form assigns the specified *wordlist* as the alias of *name; wordlist* is command and filename substituted. *Name* is not allowed to be *alias* or *unalias*

break
Causes execution to resume after the *end* of the nearest enclosing *foreach* or *while* statement. The remaining commands on the current line are executed. Multilevel breaks are thus possible by writing them all on one line.

breaksw
Causes a break from a *switch,* resuming after the *endsw.*

case label:
A label in a *switch* statement as discussed below.

cd
cd name
chdir
chdir name
Changes the shell's working directory to directory *name.* If no argument is given then changes to the home directory of the user. If *name* is not found as a subdirectory of the current directory (and does not begin with /, ./, or ../), then each

component of the variable *cdpath* is checked to see if it has a subdirectory *name*. Finally, if all else fails but *name* is a shell variable whose value begins with /, then this is tried to see if it is a directory.

continue
: Continues execution of the nearest enclosing *while* or *foreach*. The rest of the commands on the current line are executed.

default:
: Labels the default case in a *switch* statement. The default should come after all *case* labels.

echo wordlist
: The specified words are written to the shell's standard output. An \c causes the echo to complete without printing a newline. An \n in *wordlist* causes a newline to be printed. Otherwise the words are echoed, separated by spaces.

else
end
endif
endsw
: See the description of the *foreach, if, switch,* and *while* statements below.

exec command
: The specified command is executed in place of the current shell.

exit
exit(expr)
: The shell exits either with the value of the *status* variable (first form) or with the value of the specified *expr* (second form).

foreach name (wordlist)
...
end
: The variable *name* is successively set to each member of *wordlist* and the sequence of commands between this command and the matching *end* are executed. (Both *foreach* and *end* must appear alone on separate lines.)

  The built-in command *continue* may be used to continue the loop prematurely and the built-in command *break* to terminate it prematurely. When this command is read from the terminal, the loop is read up once prompting with ? before any statements in the loop are executed.

glob wordlist
: Like *echo* but no \ escapes are recognized and words are delimited by null characters in the output. Useful for programs which wish to use the shell to filename expand a list of words.

goto word
The specified *word* is filename and command expanded to yield a string of the form label. The shell rewinds its input as much as possible and searches for a line of the form label: possibly preceded by blanks or tabs. Execution continues after the specified line.

history
Displays the history event list.

if (expr) command
If the specified expression evaluates true, then the single *command* with arguments is executed. Variable substitution on *command* happens early, at the same time it does for the rest of the *if* command. *Command* must be a simple command, not a pipeline, a command list, or a parenthesized command list. Input/output redirection occurs even if *expr* is false, when command is not executed.

if (expr) **then**
...
else if (expr2) **then**
...
else
...
endif
If the specified *expr* is true then the commands to the first *else* are executed; else if *expr2* is true then the commands to the second else are executed, etc. Any number of *else-if* pairs are possible; only one *endif* is needed. The *else* part is likewise optional. (The words *else* and *endif* must appear at the beginning of input lines; the *if* must appear alone on its input line or after an $IR else .)

logout
Terminates a login shell. The only way to log out if *ignoreeof* is set.

nice
nice + number
nice command
nice + number command
The first form sets the *nice* for this shell to 4. The second form sets the *nice* to the given number. The final two forms run command at priority 4 and *number* respectively. The super-user may specify negative niceness by using "nice – number ...." The command is always executed in a subshell, and the restrictions placed on commands in simple *if* statements apply.

nohup
nohup command
The first form can be used in shell scripts to cause hangups to be ignored for the remainder of the script. The second form causes the specified command to be run with hangups ignored. Unless the shell is running detached, ***nohup*** has no effect. All processes detached with & are automatically ***nohup***ed. (Thus, ***nohup*** is not really needed.)

onintr
onintr –
onintr label
Controls the action of the shell on interrupts. The first form restores the default action of the shell on interrupts which is to terminate shell scripts or to return to the terminal command input level. The second form onintr – causes all interrupts to be ignored. The final form causes the shell to execute a goto label when an interrupt is received or a child process terminates because it was interrupted.

In any case, if the shell is running detached and interrupts are being ignored, all forms of ***onintr*** have no meaning and interrupts continue to be ignored by the shell and all invoked commands.

rehash
Causes the internal hash table of the contents of the directories in the ***path*** variable to be recomputed. This is needed if new commands are added to directories in the ***path*** while you are logged in. This should only be necessary if you add commands to one of your own directories, or if a systems programmer changes the contents of one of the system directories.

repeat count command
The specified ***command*** which is subject to the same restrictions as the ***command*** in the one line ***if*** statement above, is executed ***count*** times. I/O redirections occurs exactly once, even if ***count*** is 0.

set
set name
set name=word
set name[index]=word
set name=(wordlist)
The first form of the command shows the value of all shell variables. Variables which have other than a single word as value print as a parenthesized word list. The second form sets ***name*** to the null string. The third form sets ***name*** to the single ***word***. The fourth form sets the ***index***th component of name to word; this component must already exist. The final form sets ***name*** to the list of words in ***wordlist***. In all cases the value is command and filename expanded.

These arguments may be repeated to set multiple values in a single set command. Note however, that variable expansion happens for all arguments before any setting occurs.

setenv name value
: Sets the value of the environment variable *name* to be *value*, a single string. Useful environment variables are TERM, the type of your terminal and SHELL, the shell you are using.

shift
shift variable
: The members of *argv* are shifted to the left, discarding *argv[1]*. It is an error for *argv* not to be set or to have less than one word as value. The second form performs the same function on the specified variable.

source name
: The shell reads commands from *name*. *Source* commands may be nested; if they are nested too deeply the shell may run out of file descriptors. An error in a *source* at any level terminates all nested *source* commands. Input during *source* commands is never placed on the history list.

switch (string)
case str1:
  ...
  breaksw
...
default:
  ...
  breaksw
endsw
: Each case label is successively matched, against the specified *string* which is first command and filename expanded. The file metacharacters *, ?, and [...] may be used in the case labels, which are variable expanded. If none of the labels match before a default label is found, then the execution begins after the default label. Each case label and the default label must appear at the beginning of a line. The command *breaksw* causes execution to continue after the *endsw*. Otherwise control may fall through case labels and default labels, as in C. If no label matches and there is no default, execution continues after the *endsw*.

time
time command
: With no argument, a summary of time used by this shell and its children is printed. If arguments are given the specified simple command is timed and a time summary as described under the *time* variable is printed. If necessary, an extra shell is created to print the time statistic when the command completes.

umask
umask value
The file creation mask is displayed (first form) or set to the specified value (second form). The mask is given in octal. Common values for the mask are 002 giving all access to the group, and read and execute access to others; or 022 giving all access except no write access for users in the group or others.

unalias pattern
All aliases whose names match the specified pattern are discarded. Thus all aliases are removed by unalias *. It is not an error for nothing to be *unaliased*.

unhash
Use of the internal hash table to speed location of executed programs is disabled.

unset pattern
All variables whose names match the specified pattern are removed. Thus all variables are removed by unset *; this has noticeably distasteful side-effects. It is not an error for nothing to be *unset*.

wait
All child processes are waited for. It the shell is interactive, then an interrupt can disrupt the wait, at which time the shell prints names and process numbers of all children known to be outstanding.

while (expr)
...
end
While the specified expression evaluates nonzero, the commands between the *while* and the matching end are evaluated. *Break* and *continue* may be used to terminate or continue the loop prematurely. (The *while* and *end* must appear alone on their input lines.) Prompting occurs here the first time through the loop as for the *foreach* statement if the input is a terminal.

@
@ name = expr
@ name[index] = expr
The first form prints the values of all the shell variables. The second form sets the specified *name* to the value of *expr*. If the expression contains <, >, & or | then at least this part of the expression must be placed within ( ). The third form assigns the value of *expr* to the *index*th argument of *name*. Both *name* and its *index*th component must already exist.

Assignment operators, such as *= and + =, are available as in C. The space separating the name from the assignment operator is optional. Spaces are mandatory in separating components of *expr* which would otherwise be single words.

Special postfix + + and – – operators increment and decrement *name* respectively, i.e. @ i+ + .

Predefined Variables

The following variables have special meaning to the shell. Of these, *argv, child, home, path, prompt, shell* and *status* are always set by the shell. Except for *child* and *status* this setting occurs only at initialization; these variables will not then be modified unless done explicitly by the user.

The shell copies the environment variable PATH into the variable *path*, and copies the value back into the environment whenever *path* is set. Thus is is not necessary to worry about its setting other than in the file *.cshrc* as inferior *csh* processes will import the definition of *path* from the environment.

argv Set to the arguments to the shell, it is from this variable that positional parameters are substituted, i.e. **$1** is replaced by **$argv[1]**, etc.

cdpath Gives a list of alternate directories searched to find subdirectories in *cd* commands.

child The process number printed when the last command was forked with &. This variable is *unset* when this process terminates.

echo Set when the – x command line option is given. Causes each command and its arguments to be echoed just before it is executed. For nonbuilt-in commands all expansions occur before echoing. Built-in commands are echoed before command and filename substitution, since these substitutions are then done selectively.

histchars Can be assigned a two-character string. The first character is used as a history character in place of !, the second character is used in place of the ^ substitution mechanism. For example, set histchars=",;" will cause the history characters to be comma and semicolon.

history Can be given a numeric value to control the size of the history list. Any command which has been referenced in this many events will not be discarded. A *history* that is too large may run the shell out of memory. The last executed command is always saved on the history list.

home The home directory of the invoker, initialized from the environment. The filename expansion of ~ refers to this variable.

ignoreeof — If set the shell ignores end-of-file from input devices that are terminals. This prevents a shell from accidentally being terminated by typing a CNTRL-D.

mail — The files where the shell checks for mail. This is done after each command completion which will result in a prompt, if a specified interval has elapsed. The shell says "You have new mail", if the file exists with an access time not greater than its modify time.

If the first word of the value of *mail* is numeric it specifies a different mail checking interval, in seconds, than the default, which is 10 minutes.

If multiple mail files are specified, then the shell says "New mail in *name*" when there is mail in the file *name*.

noclobber — As described in the section *Input/output,* restrictions are placed on output redirection to insure that files are not accidentally destroyed, and that >> redirections refer to existing files.

noglob — If set, filename expansion is inhibited. This is most useful in shell scripts which are not dealing with filenames, or after a list of filenames has been obtained and further expansions are not desirable.

nonomatch — If set, it is not an error for a filename expansion to not match any existing files; rather, the primitive pattern is returned. It is still an error for the primitive pattern to be malformed, i.e. echo [still gives an error.

path — Each word of the path variable specifies a directory in which commands are to be sought for execution. A null word specifies the current directory. If there is no *path* variable then only full pathnames will execute. The usual search path is /bin, /usr/bin, and ., but this may vary from system to system. For the super-user the default search path is /etc, /bin and /usr/bin. A shell which is given neither the – c nor the – t option will normally hash the contents of the directories in the *path* variable after reading *.cshrc,* and each time the *path* variable is reset. If new commands are added to these directories while the shell is active, it may be necessary to give the *rehash* or the commands may not be found.

prompt — The string which is printed before each command is read from an interactive terminal input. If a !

appears in the string it will be replaced by the current event number unless a preceding \ is given. Default is %, or # for the super-user.

shell The file in which the shell resides. This is used in forking shells to interpret files which have execute bits set, but which are not executable by the system. (See the description of *Nonbuilt-In Command Execution* below.) Initialized to the (system-dependent) home of the shell.

status The status returned by the last command. If it terminated abnormally, then **0200** is added to the status. Abnormal termination results in a core dump. Built-in commands which fail return exit status **1**, all other built-in commands set status **0**.

time Controls automatic timing of commands. If set, then any command which takes more than this many cpu seconds will cause a line giving user, system, and real times and a utilization percentage which is the ratio of user plus system times to real time to be printed when it terminates.

verbose Set by the – v command line option, causes the words of each command to be printed after history substitution.

Nonbuilt-In Command Execution

When a command to be executed is found to not be a built-in command the shell attempts to execute the command via *exec*(S). Each word in the variable *path* names a directory from which the shell will attempt to execute the command. If it is given neither a – c nor a – t option, the shell will hash the names in these directories into an internal table so that it will only try an *exec* in a directory if there is a possibility that the command resides there. This greatly speeds command location when a large number of directories are present in the search path. If this mechanism has been turned off (via *unhash),* or if the shell was given a – c or – t argument, and in any case for each directory component of *path* which does not begin with a /, the shell concatenates with the given command name to form a pathname of a file which it then attempts to execute.

Parenthesized commands are always executed in a subshell. Thus (cd ; pwd) ; pwd prints the *home* directory; leaving you where you were (printing this after the home directory), while cd ; pwd leaves you in the home directory. Parenthesized commands are most often used to prevent *cd* from affecting the current shell.

If the file has execute permissions but is not an executable binary to the system, then it is assumed to be a file containing shell commands and a new shell is spawned to read it.

If there is an *alias* for *shell* then the words of the alias will be prepended to the argument list to form the shell command. The first word of the *alias* should be the full pathname of the shell (e.g. $shell). Note that this is a special, late occurring, case of *alias* substitution, and only allows words to be prepended to the argument list without modification.

Argument List Processing

If argument 0 to the shell is - , then this is a login shell. The flag arguments are interpreted as follows:

- c Commands are read from the (single) following argument which must be present. Any remaining arguments are placed in *argv*.

- e The shell exits if any invoked command terminates abnormally or yields a nonzero exit status.

- f The shell will start faster, because it will neither search for nor execute commands from the file .cshrc in the invoker's home directory.

- i The shell is interactive and prompts for its top-level input, even if it appears to not be a terminal. Shells are interactive without this option if their inputs and outputs are terminals.

- n Commands are parsed, but not executed. This may aid in syntactic checking of shell scripts.

- s Command input is taken from the standard input.

- t A single line of input is read and executed. A \ may be used to escape the newline at the end of this line and continue onto another line.

- v Causes the *verbose* variable to be set, with the effect that command input is echoed after history substitution.

- x Causes the *echo* variable to be set, so that commands are echoed immediately before execution.

- V Causes the *verbose* variable to be set even before .cshrc is executed.

- X Causes the *echo* variable to be set even before .cshrc is executed.

After processing of flag arguments, if arguments remain but none of the - c, - i, - s, or - t options were given, the first argument is taken as the name of a file of commands to be executed. The shell opens this file, and saves its name for possible resubstitution by $0. Since on a typical system most shell scripts are written for the

standard shell (see *sh*(C)), the C shell will execute such a standard shell if the first character of a script is not a #, i.e. if the script does not start with a comment. Remaining arguments initialize the variable *argv*.

Signal Handling

The shell normally ignores *quit* signals. The *interrupt* and *quit* signals are ignored for an invoked command if the command is followed by &; otherwise the signals have the values which the shell inherited from its parent. The shells handling of interrupts can be controlled by *onintr*. Login shells catch the *terminate* signal; otherwise this signal is passed on to children from the state in the shell's parent. In no case are interrupts allowed when a login shell is reading the file .logout.

Files

~/.cshrc	Read at by each shell at the beginning of execution
~/.login	Read by login shell, after .cshrc at login
~/.logout	Read by login shell, at logout
/bin/sh	Shell for scripts not starting with a #
/tmp/sh*	Temporary file for <<
/dev/null	Source of empty file
/etc/passwd	Source of home directories for ~name

Limitations

Words can be no longer than 512 characters. The number of arguments to a command which involves filename expansion is limited to 1/6 number of characters allowed in an argument list, which is 5120, less the characters in the environment. Also, command substitutions may substitute no more characters than are allowed in an argument list.

To detect looping, the shell restricts the number of *aliae* substitutions on a single line to 20.

See Also

access(S), exec(S), fork(S), pipe(S), signal(S), umask(S), wait(S), a.out(F), environ(M)

Credit

This utility was developed at the University of California at Berkeley and is used with permission.

Notes

Built-in control structure commands like **foreach** and **while** cannot be used with | & or ;.

Commands within loops, prompted for by ?, are not placed in the *history* list.

It is not possible to use the colon (:) modifiers on the output of command substitutions.

Csh attempts to import and export the PATH variable for use with regular shell scripts. This only works for simple cases, where the PATH contains no command characters.

This version of *csh* does not support or use the process control features of the 4th Berkeley Distribution.

Name

ctags – Creates a tags file.

Syntax

ctags [– u] [– w] [– x] name ...

Description

Ctags makes a tags file for *vi*(C) from the specified C sources. A tags file gives the locations of specified objects (in this case functions) in a group of files. Each line of the tags file contains the function name, the file in which it is defined, and a scanning pattern used to find the function definition. These are given in separate fields on the line, separated by blanks or tabs. Using the *tags* file, *vi* can quickly find these function definitions.

If the – x flag is given, *ctags* produces a list of function names, the line number and file name on which each is defined, as well as the text of that line and prints this on the standard output. With the – x option no tags file is created. This is a simple index which can be printed out as an off-line readable function index.

Files whose name ends in .c or .h are assumed to be C source files and are searched for C routine and macro definitions.

Other options are:

– w Suppresses warning diagnostics.

– u Causes the specified files to be *updated* in tags; that is, all references to them are deleted, and the new values are appended to the file. (Beware: this option is implemented in a way which is rather slow; it is usually faster to simply rebuild the *tags* file.)

The tag *main* is treated specially in C programs. The tag formed is created by prepending *M* to the name of the file, with a trailing .c removed, if any, and leading pathname components also removed. This makes use of *ctags* practical in directories with more than one program.

Files

tags	Output tags file

See Also

ex(C), vi(C)

Credit

This utility was developed at the University of California at Berkeley and is used with permission.

Name

delta – Makes a delta (change) to an SCCS file.

Syntax

delta [– rSID] [– s] [– n] [– glist] [– m[mrlist]] [– y[comment]] [– p] files

Description

Delta is used to permanently introduce into the named SCCS file changes that were made to the file retrieved by *get*(CP) (called the *g-file*, or generated file).

Delta makes a delta to each SCCS file named by *files*. If a directory is named, *delta* behaves as though each file in the directory were specified as a named file, except that nonSCCS files (last component of the pathname does not begin with s.) and unreadable files are silently ignored. If a name of – is given, the standard input is read (see Warning); each line of the standard input is taken to be the name of an SCCS file to be processed.

Delta may issue prompts on the standard output depending upon certain options specified and flags (see *admin*(CP)) that may be present in the SCCS file (see – m and – y options below).

Options apply independently to each named file.

– r*SID* Uniquely identifies which delta is to be made to the SCCS file. The use of this keyletter is necessary only if two or more versions of the same SCCS file have been retrieved for editing (get – e) by the same person (login name). The SID value specified with the – r keyletter can be either the SID specified on the *get* command line or the SID to be made as reported by the *get* command (see *get*(CP)). A diagnostic results if the specified SID is ambiguous, or if it is necessary and omitted on the command line.

– s Suppresses the issue, on the standard output, of the created delta's SID, as well as the number of lines inserted, deleted and unchanged in the SCCS file.

– n Specifies retention of the edited *g-file* (normally removed at completion of delta processing).

– g*list* Specifies a *list* (see *get*(CP) for the definition of *list*) of deltas which are to be *ignored* when the file is accessed at the change level (SID) created by this delta.

– m[*mrlist*] If the SCCS file has the v flag set (see *admin*(CP)) then a Modification Request (MR) number *must* be supplied as the reason for creating the new delta.

If – m is not used and the standard input is a terminal, the prompt **MRs?** is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. The **MRs?** prompt always precedes the **comments?** prompt (see – y keyletter).

MRs in a list are separated by blanks and/or tab characters. An unescaped newline character terminates the MR list.

Note that if the v flag has a value (see *admin*(CP)), it is taken to be the name of a program (or shell procedure) which will validate the correctness of the MR numbers. If a nonzero exit status is returned from MR number validation program, *delta* terminates (it is assumed that the MR numbers were not all valid).

– y[*comment*] Arbitrary text used to describe the reason for making the delta. A null string is considered a valid *comment*.

If – y is not specified and the standard input is a terminal, the prompt **comments?** is issued on the standard output before the standard input is read; if the standard input is not a terminal, no prompt is issued. An unescaped newline character terminates the comment text.

– p Causes *delta* to print (on the standard output) the SCCS file differences before and after the delta is applied. Differences are displayed in a *diff*(C) format.

Files

All files of the form *?*-file are explained in Chapter 5, "SCCS: A Source Code Control System" in the *XENIX Programmer's Guide*. The naming convention for these files is also described there.

g-file Existed before the execution of *delta*; removed after completion of *delta*.

p-file	Existed before the execution of *delta*; may exist after completion of *delta*.
q-file	Created during the execution of *delta*; removed after completion of *delta*.
x-file	Created during the execution of *delta*; renamed to SCCS file after completion of *delta*.
z-file	Created during the execution of *delta*; removed during the execution of *delta*.
d-file	Created during the execution of *delta*; removed after completion of *delta*.
/usr/bin/bdiff	Program to compute differences between the "retrieved" file and the *g-file*.

Warning

Lines beginning with an **SOH** ASCII character (binary **001**) cannot be placed in the SCCS file unless the **SOH** is escaped. This character has special meaning to SCCS (see *sccsfile*(F)) and will cause an error.

A *get* of many SCCS files, followed by a *delta* of those files, should be avoided when the *get* generates a large amount of data. Instead, multiple *get/delta* sequences should be used.

If the standard input (–) is specified on the *delta* command line, the – **m** (if necessary) and – **y** options *must* also be present. Omission of these options causes an error to occur.

See Also

admin(CP), bdiff(C), get(CP), help(CP), prs(CP), sccsfile(F)

Diagnostics

Use *help*(CP) for explanations.

Name

get – Gets a version of an SCCS file.

Syntax

get [– **r**SID] [– **c**cutoff] [– **i**list] [– **x**list] [– aseq-no.] [– **k**] [– **e**] [– **l**[p]] [– **p**] [– **m**] [– **n**] [– **s**] [– **b**] [– **g**] [– **t**] file ...

Description

Get generates an ASCII text file from each named SCCS file according to the specifications given by its options, which begin with – . The arguments may be specified in any order, but all options apply to all named SCCS files. If a directory is named, *get* behaves as though each file in the directory were specified as a named file, except that nonSCCS files (last component of the pathname does not begin with s.) and unreadable files are silently ignored. If a name of – is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed. Again, nonSCCS files and unreadable files are silently ignored.

The generated text is normally written into a file called the *g-file* whose name is derived from the SCCS filename by simply removing the leading **s.**; (see also *FILES*, below).

Each of the options is explained below as though only one SCCS file is to be processed, but the effects of any option apply independently to each named file.

– **r***SID* The *S*CCS *ID*entification string (SID) of the version (delta) of an SCCS file to be retrieved.

– **c***cutoff* *Cutoff* date-time, in the form:

YY[MM[DD[HH[MM[SS]]]]]

No changes (deltas) to the SCCS file that were created after the specified *cutoff* date-time are included in the generated ASCII text file. Units omitted from the date-time default to their maximum possible values; that is, – **c7502** is equivalent to – **c750228235959**. Any number of nonnumeric characters may separate the various 2 digit pieces of the *cutoff* date-time. This feature allows you to specify a *cutoff* date in the form: "– c77/2/2 9:22:25".

– **e** Indicates that the *get* is for the purpose of editing or making a change (delta) to the SCCS file via a subsequent use of *delta*(CP). The – **e** option used in a *get* for a particular version (SID) of the SCCS file prevents further

gets for editing on the same SID until *delta* is executed or the **j** (joint edit) flag is set in the SCCS file (see *admin*(CP)). Concurrent use of **get** – **e** for different SIDs is always allowed.

If the *g-file* generated by *get* with an – **e** option is accidentally ruined in the editing process, it may be regenerated by reexecuting the *get* command with the – **k** option in place of the – **e** option.

SCCS file protection specified via the ceiling, floor, and authorized user list stored in the SCCS file (see *admin*(CP)) are enforced when the – **e** option is used.

– **b** Used with the – **e** option to indicate that the new delta should have an SID in a new branch. This option is ignored if the **b** flag is not present in the file (see *admin*(CP)) or if the retrieved *delta* is not a leaf *delta*. (A leaf *delta* is one that has no successors on the SCCS file tree.)

Note: A branch *delta* may always be created from a non-leaf *delta*.

– **i***list* A *list* of deltas to be included (forced to be applied) in the creation of the generated file. The *list* has the following syntax:

<list> ::= <range> | <list> , <range>
<range> ::= SID | SID – SID

SID, the SCCS Identification of a delta, may be in any form described in Chapter 5, "SCCS: A Source Code Control System," in the XENIX *Programmer's Guide*.

– **x***list* A *list* of deltas to be excluded (forced not to be applied) in the creation of the generated file. See the – **i** option for the *list* format.

– **k** Suppresses replacement of identification keywords (see below) in the retrieved text by their value. The – **k** option is implied by the – **e** option.

– **l**[**p**] Causes a delta summary to be written into an *l-file*. If – **lp** is used then an *l-file* is not created; the delta summary is written on the standard output instead. See *FILES* for the format of the *l-file*.

– **p** Causes the text retrieved from the SCCS file to be written on the standard output. No *g-file* is created. All output that normally goes to the standard output goes to file descriptor 2 instead, unless the – **s** option is used, in which case it disappears.

– **s** Suppresses all output normally written on the standard output. However, fatal error messages (which always go to file descriptor 2) remain unaffected.

– **m** Causes each text line retrieved from the SCCS file to be preceded by the SID of the delta that inserted the text line in the SCCS file. The format is: SID, followed by a horizontal tab, followed by the text line.

– **n** Causes each generated text line to be preceded with the %M% identification keyword value (see below). The format is: %M% value, followed by a horizontal tab, followed by the text line. When both the – **m** and – **n** options are used, the format is: %M% value, followed by a horizontal tab, followed by the – **m** option generated format.

– **g** Suppresses the actual retrieval of text from the SCCS file. It is primarily used to generate an *l-file*, or to verify the existence of a particular SID.

– **t** Used to access the most recently created (top) delta in a given release (e.g., – **r1**), or release and level (e.g., – **r1.2**).

– **a***seq-no.* The delta sequence number of the SCCS file delta (version) to be retrieved (see *sccsfile*(F)). This option is used by the *comb*(CP) command; it is not particularly useful should be avoided. If both the – **r** and – **a** options are specified, the – **a** option is used. Care should be taken when using the – **a** option in conjunction with the – **e** option, as the SID of the delta to be created may not be what you expect. The – **r** option can be used with the – **a** and – **e** options to control the naming of the SID of the delta to be created.

For each file processed, *get* responds (on the standard output) with the SID being accessed and with the number of lines retrieved from the SCCS file.

If the – **e** option is used, the SID of the delta to be made appears after the SID accessed and before the number of lines generated. If there is more than one named file or if a directory or standard input is named, each filename is printed (preceded by a newline) before it is processed. If the – **i** option is used included deltas are listed following the notation "Included"; if the – **x** option is used, excluded deltas are listed following the notation "Excluded".

Identification Keywords

Identifying information is inserted into the text retrieved from the SCCS file by replacing *identification keywords* with their value

wherever they occur. The following keywords may be used in the text stored in an SCCS file:

Keyword	Value
%M%	Module name: either the value of the m flag in the file (see *admin*(CP)), or if absent, the name of the SCCS file with the leading s. removed.
%I%	SCCS identification (SID) (%R%.%L%.%B%.%S%) of the retrieved text.
%R%	Release.
%L%	Level.
%B%	Branch.
%S%	Sequence.
%D%	Current date (YY/MM/DD).
%H%	Current date (MM/DD/YY).
%T%	Current time (HH:MM:SS).
%E%	Date newest applied delta was created (YY/MM/DD).
%G%	Date newest applied delta was created (MM/DD/YY).
%U%	Time newest applied delta was created (HH:MM:SS).
%Y%	Module type: value of the t flag in the SCCS file (see *admin*(CP)).
%F%	SCCS filename.
%P%	Fully qualified SCCS filename.
%Q%	The value of the q flag in the file (see *admin*(CP)).
%C%	Current line number. This keyword is intended for identifying messages output by the program such as "this shouldn't have happened" type errors. It is *not* intended to be used on every line to provide sequence numbers.
%Z%	The 4-character string @ (#) recognizable by *what*(C).
%W%	A shorthand notation for constructing *what*(C) strings for XENIX program files. %W% = %Z%%M%<horizontal-tab>%I%
%A%	Another shorthand notation for constructing *what*(C) strings for nonXENIX program files. %A% = %Z%%Y% %M% %I%%Z%

Files

Several auxiliary files may be created by *get*. These files are known generically as the *g-file*, *l-file*, *p-file*, and *z-file*. The letter before the hyphen is called the tag. An auxiliary filename is formed from the SCCS filename: the last component of all SCCS filenames must be of the form s.*module-name*, the auxiliary files are named by replacing the leading s with the tag. The *g-file* is an exception to this scheme: the *g-file* is named by removing the s. prefix. For example, s.xyz.c, the auxiliary filenames would be xyz.c, l.xyz.c, p.xyz.c, and z.xyz.c, respectively.

The *g-file*, which contains the generated text, is created in the current directory (unless the – p option is used). A *g-file* is created in all cases, whether or not any lines of text were generated by the *get*. It is owned by the real user. If the – k option is used or

implied, the *g-file*'s mode is 644; otherwise the mode is 444. Only the real user need have write permission in the current directory.

The *l-file* contains a table showing which deltas were applied in generating the retrieved text. The *l-file* is created in the current directory if the – l option is used; its mode is 444 and it is owned by the real user. Only the real user need have write permission in the current directory.

Lines in the *l-file* have the following format:

a. A blank character if the delta was applied;
   * otherwise
b. A blank character if the delta was applied or wasn't applied and ignored;
   * if the delta wasn't applied and wasn't ignored
c. A code indicating a "special" reason why the delta was or was not applied:
   "I": Included
   "X": Excluded
   "C": Cut off (by a – c option)
d. Blank
e. SCCS identification (SID)
f. Tab character
g. Date and time (in the form YY/MM/DD HH:MM:SS) of creation
h. Blank
i. Login name of person who created *delta*

The comments and MR data follow on subsequent lines, indented one horizontal tab character. A blank line terminates each entry.

The *p-file* is used to pass information resulting from a *get* with an – e option along to *delta*. Its contents are also used to prevent a subsequent execution of *get* with an – e option for the same SID until *delta* is executed or the joint edit flag, j, (see *admin*(CP)) is set in the SCCS file. The *p-file* is created in the directory containing the SCCS file and the effective user must have write permission in that directory. Its mode is 644 and it is owned by the effective user. The format of the *p-file* is: the gotten SID, followed by a blank, followed by the SID that the new delta will have when it is made, followed by a blank, followed by the login name of the real user, followed by a blank, followed by the date-time the *get* was executed, followed by a blank and the – i option if it was present, followed by a blank and the – x option if it was present, followed by a newline. There can be an arbitrary number of lines in the *p-file* at any time; no two lines can have the same new delta SID.

The *z-file* serves as a *lock-out* mechanism against simultaneous updates. Its contents are the binary (2 bytes) process ID of the command (i.e., *get*) that created it. The *z-file* is created in the directory containing the SCCS file for the duration of *get*. The same protection restrictions as those for the *p-file* apply for the *z-file*. The *z-file* is

created mode 444.

See Also

admin(CP), delta(CP), help(CP), prs(CP), what(C), sccsfile(F)

Diagnostics

Use *help*(CP) for explanations.

Notes

If the effective user has write permission (either explicitly or implicitly) in the directory containing the SCCS files, but the real user doesn't, then only one file may be named when the **– e** option is used.

Name

gets – Gets a string from the standard input.

Syntax

gets [string]

Description

Gets can be used with ***csh***(CP) to read a string from the standard input. If ***string*** is given it is used as a default value if an error occurs. The resulting string (either ***string*** or as read from the standard input) is written to the standard output. If no ***string*** is given and an error occurs, ***gets*** exits with exit status 1.

See Also

line(C), csh(CP)

Name

hdr – Displays selected parts of object files.

Syntax

hdr [– dhprsSt] file ...

Description

Hdr displays object file headers, symbol tables, and text or data relocation records in human-readable formats. It also prints out seek positions for the various segments in the object file.

A.out, x.out, and x.out segmented formats and archives are understood.

The symbol table format consists of six fields. In **a.out** formats the third field is missing. The first field is the symbol's index or position in the symbol table, printed in decimal. The index of the first entry is zero. The second field is the type, printed in hexadecimal. The third field is the **s_seg** field, printed in hexadecimal. The fourth field is the symbol's value in hexadecimal. The fifth field is a single character which represents the symbol's type as in *nm*(CP), except C common is not recognized as a special case of undefined. The last field is the symbol name.

If long form relocation is present, the format consists of six fields. The first is the descriptor, printed in hexadecimal. The second is the symbol ID, or index, in decimal. This field is used for external relocations as an index into the symbol table. It should reference an undefined symbol table entry. The third field is the position, or offset, within the current segment at which relocation is to take place; it is printed in hexadecimal. The fourth field is the name of the segment referenced in the relocation: text, data, bss or EXT for external. The fifth field is the size of relocation: byte, word (2 bytes), or long. The last field will indicate, if present, that the relocation is relative.

If short form relocation is present, the format consist of three fields. The first field is the relocation command in hexadecimal. the second field contains the name of the segment referenced; text or data. The last field indicates the size of relocation: word or long.

Options and their meanings are:

– **h** Causes the object file header and extended header to be printed out. Each field in the header or extended header is labeled. This is the default option.

– **d** Causes the data relocation records to be printed out.

– **t** Causes the text relocation records to be printed out.

– **r** Causes both text and data relocation to be printed.

– **p** Causes seek positions to be printed out as defined by macros in the include file, **<a.out.h>**.

– **s** Prints the symbol table.

– **S** Prints the file segment table with a header. (Only applicable to **x.out** segmented executable files.)

See Also

a.out(F), nm(CP)

Name

help – Asks for help about SCCS commands.

Syntax

help [args]

Description

Help finds information to explain a message from an SCCS command or explain the use of a command. Zero or more arguments may be supplied. If no arguments are given, *help* will prompt for one.

The arguments may be either message numbers (which normally appear in parentheses following messages) or command names. There are the following types of arguments:

type 1 Begins with nonnumerics, ends in numerics. The nonnumeric prefix is usually an abbreviation for the program or set of routines which produced the message (e.g., **ge6**, for message 6 from the *get* command).

type 2 Does not contain numerics (as a command, such as **get**)

type 3 Is all numeric (e.g., **212**)

The response of the program will be the explanatory information related to the argument, if there is any.

When all else fails, try "help stuck".

Files

/usr/lib/help Directory containing files of message text

Name

ld – Invokes the link editor.

Syntax

ld [*options*] *filename...*

Description

Ld is the XENIX link editor. It creates an executable program by combining one or more object files and copying the executable result to the file **a.out**. The *filename* must name an object or library file. These names must have the ".o" (for object) or ".a" (for archive library) extensions. If more than one name is given, the names must be separated by one or more spaces. If errors occur while linking, *ld* displays an error message; the resulting **a.out** file is unexecutable.

Ld concatenates the contents of the given object files in the order given in the command line. Library files in the command line are examined only if there are unresolved external references encountered from previous object files. Library files must be in *ranlib*(CP) format, that is, the first member must be named __.SYMDEF, which is a dictionary for the library. *Ld* ignores the modification dates of the library and the __.SYMDEF entry, so if object files have been added to the library since __.SYMDEF was created, the link may result in an "invalid object module."

The library is searched iteratively to satisfy as many references as possible and only those routines that define unresolved external references are concatenated. Object and library files are processed at the point they are encountered in the argument list, so the order of files in the command line is important. In general, all object files should be given before library files. *Ld* sets the entry point of the resulting program to the beginning of the first routine.

There are the following options:

– A*num*
Creates a standalone program whose expected load address (in hexadecimal) is *num*. This option sets the absolute flag in the header of the a.out file. Such program files can only be executed as standalone programs.

– F*num*
Sets the size of the program stack to *num* bytes. Default stack size if not given, is 2 Kbytes.

– i Creates separate instruction and data spaces for small model programs. When the output file is executed, the program text

and data areas are allocated separate physical segments. The text portion will be read-only and shared by all users executing the file.

– Ms
Creates small model program and checks for error, such as fixup overflow. This option is reserved for object files compiled or assembled using the small model configuration. This is the default model if no – M option is given.

– Mm
Creates middle model program and checks for errors. This option is reserved for object files compiled or assembled using the middle model configuration. This option implies – i .

– Ml
Creates a large model program and checks for errors. The option is reserved for object files compiled using the large model configuration. This option implies – i .

– o *name*
Sets the executable program filename to *name* instead of a.out.

Ld should be invoked using the *cc*(CP) instead of invoking it directly. *Cc* invokes *ld* as the last step of compilation, providing all the necessary C-language support routines. Invoking *ld* directly is not recommended since failure to give command line arguments in the correct order can result in errors.

Files

/bin/ld

See Also

as(CP), ar(CP), cc(CP), ranlib(CP)

Notes

The user must make sure that the most recent library versions have been processed with *ranlib*(CP) before linking. If this is not done, *ld* cannot create executable programs using these libraries.

Name

lex – Generates programs for lexical analysis.

Syntax

lex [– ctvn] [file] ...

Description

Lex generates programs to be used in simple lexical analysis of text. A file lex.yy.c is generated which, when loaded with the *lex* library, copies the input to the output except when a string specified in the file is found. If a string is found, the corresponding program text is executed.

The input *file* contains strings and expressions to be searched for, and C text to be executed when strings are found. Multiple files are treated as a single file. If no files are specified, standard input is used.

The options must appear before any files. The options are as follows:

– c Indicates C actions and is the default.

– t Causes the lex.yy.c program to be written instead to standard output.

– v Provides a one-line summary of statistics of the machine generated.

– n Suppresses the – summary.

Strings and Operators

Lex strings may contain square brackets to indicate character classes, as in [abx– z] to indicate a, b, x, y, and z; and the operators *, +, and ? mean respectively any nonnegative number of, any positive number of, and either zero or one occurrences of, the previous character or character class. Thus, [a-zA-Z]+ matches a string of letters. The character . is the class of all ASCII characters except newline. Parentheses for grouping and vertical bar for alternation are also supported. The notation *r*{*d*,*e*} in a rule indicates between *d* and *e* instances of regular expression *r*. It has higher precedence than | but lower than *, ?, +, and concatenation. The character ^ at the beginning of an expression permits a successful match only immediately after a newline, and the character $ at the end of an expression requires a trailing newline. The character / in an expression indicates trailing context; only the part of the expression up to

the slash is returned in *yytext*, but the remainder of the expression must follow in the input stream. An operator character may be used as an ordinary symbol if it is within " symbols or preceded by \.

Routines and Variables

Matching is done in order of the strings in the file. The actual string matched is left in *yytext*, an external character array. Three subroutines defined as macros are expected: **input()** to read a character; **unput(***c***)** to replace a character read; and **output(***c***)** to place an output character. They are defined in terms of the standard streams, but you can override them. The program generated is named yylex(), and the library contains a **main()** which calls it. The action REJECT on the right side of the rule causes this match to be rejected and the next suitable match executed; the function **yymore()** accumulates additional characters into the same *yytext*; and the function yyless(*p*) pushes back the portion of the string matched beginning at *p*, which should be between *yytext* and *yytext*+ *yyleng*. The macros *input* and *output* use files yyin and yyout to read from and write to, defaulted to stdin and stdout, respectively. The external names generated by *lex* all begin with the prefix yy or YY.

Lex File Format

Any line beginning with a blank is assumed to contain only C text and is copied; if it precedes %% it is copied into the external definition area of the lex.yy.c file. All rules should follow a %%, as in YACC. Lines preceding %% which begin with a nonblank character define the string on the left to be the remainder of the line; it can be called out later by surrounding it with {}. Note that curly brackets do not imply parentheses; only string substitution is done.

Certain table sizes for the resulting finite state machine can be set in the definitions section:

%p *n*
number of positions is *n* (default 2000)

%n *n*
number of states is *n* (500)

%t *n*
number of parse tree nodes is *n* (1000)

%a *n*
number of transitions is *n* (3000)

The use of one or more of the above automatically implies the **– v** option, unless the **– n** option is used.

Example

```
D       [0-9]
%%
if      printf("IF statement\n");
[a-z]+  printf("tag, value %s\n",yytext);
0{D}+   printf("octal number %s\n",yytext);
{D}+    printf("decimal number %s\n",yytext);
"++"    printf("unary op\n");
"+"     printf("binary op\n");
"/*"    {       loop:
                while (input() != '*');
                switch (input())
                        {
                        case '/': break;
                        case '*': unput('*');
                        default: go to loop;
                        }
                }
```

See Also

yacc(CP)
XENIX *Programmer's Guide*

Notes

This program translates its input into C source code, which in segmented programming environments, is suitable for compiling as a small model program only (see *cc*(CP)).

Name

lint – Checks C language usage and syntax.

Syntax

lint [– **abchlnpuvx**] file ...

Description

Lint attempts to detect features of the C program *file* that are likely to be bugs, nonportable, or wasteful. It also checks type usage more strictly than the C compiler. Among the things which are currently detected are unreachable statements, loops not entered at the top, automatic variables declared and not used, and logical expressions whose value is constant. Moreover, the usage of functions is checked to find functions which return values in some places and not in others, functions called with varying numbers of arguments, and functions whose values are not used.

If more than one *file* is given, it is assumed that all the files are to be loaded together; they are checked for mutual compatibility. If routines from the standard library are called from *file*, *lint* checks the function definitions using the standard lint library **llibc.ln**. If *lint* is invoked with the – p option, it checks function definitions from the portable lint library **llibport.ln**.

Any number of *lint* options may be used, in any order. The following options are used to suppress certain kinds of complaints:

– a Suppresses complaints about assignments of long values to variables that are not long.

– b Suppresses complaints about **break** statements that cannot be reached. (Programs produced by *lex* or *yacc* will often result in a large number of such complaints.)

– c Suppresses complaints about casts that have questionable portability.

– h Does not apply heuristic tests that attempt to intuit bugs, improve style, and reduce waste.

– u Suppresses complaints about functions and external variables used and not defined, or defined and not used. (This option is suitable for running *lint* on a subset of files of a larger program.)

– v Suppresses complaints about unused arguments in functions.

– x Does not report variables referred to by external declarations but never used.

The following arguments alter *lint's* behavior:

– **n** Does not check compatibility against either the standard or the portable lint library.

– **p** Attempts to check portability to other dialects of C.

– llibname
Checks functions definitions in the specified lint library. For example, – **lm** causes the library *llibm.ln* to be checked.

The – **D**, – **U**, and – **I** options of *cc*(CP) are also recognized as separate arguments.

Certain conventional comments in the C source will change the behavior of *lint*:

/*NOTREACHED*/
At appropriate points stops comments about unreachable code.

/*VARARGS*n*/
Suppresses the usual checking for variable numbers of arguments in the following function declaration. The data types of the first *n* arguments are checked; a missing *n* is taken to be 0.

/*ARGSUSED*/
Turns on the – **v** option for the next function. .

/*LINTLIBRARY*/
Shuts off complaints about unused functions in this file.

Lint produces its first output on a per source file basis. Complaints regarding included files are collected and printed after all source files have been processed. Finally, information gathered from all input files is collected and checked for consistency. At this point, if it is not clear whether a complaint stems from a given source file or from one of its included files, the source filename will be printed followed by a question mark.

Files

/usr/lib/lint[12] Program files

/usr/lib/llibc.ln, /usr/lib/llibport.ln, /usr/lib/llibm.ln, /usr/lib/llibdbm.ln, /usr/lib/llibtermlib.ln
Standard lint libraries (binary format)

/usr/lib/llibc, /usr/lib/llibport, /usr/lib/llibm, /usr/lib/llibdbm, /usr/lib/llibtermlib
Standard lint libraries (source format)

/usr/tmp/*lint* Temporaries

See Also

cc(CP)

Notes

Exit(S), and other functions which do not return, are not understood. This can cause improper error messages.

Name

lorder – Finds ordering relation for an object library.

Syntax

lorder file ...

Description

Lorder creates an ordered listing of object filenames, showing which files depend on variables declared in other files. The *file* is one or more object or library archive files (see *ar*(CP)). The standard output is a list of pairs of object filenames. The first file of the pair refers to external identifiers defined in the second. The output may be processed by *tsort*(CP) to find an ordering of a library suitable for one-pass access by *ld*(CP).

Example

The following command builds a new library from existing .o files:

```
ar cr library `lorder *.o |tsort`
```

Files

*symref, *symdef Temp files

See Also

ar(CP), ld(CP), tsort(CP)

Notes

Object files whose names do not end with .o, even when contained in library archives, are overlooked. Their global symbols and references are attributed to some other file.

Name

m4 – Invokes a macro processor.

Syntax

m4 [options] [files]

Description

M4 is a macro processor intended as a front end for Ratfor, C, and other languages. Each of the argument *files* is processed in order; if there are no files, or if a filename is – , the standard input is read. The processed text is written on the standard output.

The options and their effects are as follows:

– **e** Operates interactively. Interrupts are ignored and the output is unbuffered.

– **s** Enables line sync output for the C preprocessor (**#**line ...)

– **B***int*
Changes the size of the push-back and argument collection buffers from the default of 4,096.

– **H***int*
Changes the size of the symbol table hash array from the default of 199. The size should be prime.

– **S***int*
Changes the size of the call stack from the default of 100 slots. Macros take three slots, and nonmacro arguments take one.

– **T***int*
Changes the size of the token buffer from the default of 512 bytes.

To be effective, these flags must appear before any filenames and before any – **D** or – **U** flags:

– **D***name*[=*val*]
Defines *name* to *val* or to null in *val*'s absence.

– **U***name*
Undefines *name*.

Macro Calls

Macro calls have the form:

name(arg1,arg2, ..., argn)

The (must immediately follow the name of the macro. If a defined macro name is not followed by a (, it is deemed to have no arguments. Leading unquoted blanks, tabs, and newlines are ignored while collecting arguments. Potential macro names consist of alphabetic letters, digits, and underscore _, where the first character is not a digit.

Left and right single quotation marks are used to quote strings. The value of a quoted string is the string stripped of the quotation marks.

When a macro name is recognized, its arguments are collected by searching for a matching right parenthesis. Macro evaluation proceeds normally during the collection of the arguments, and any commas or right parentheses which happen to turn up within the value of a nested call are as effective as those in the original input text. After argument collection, the value of the macro is pushed back onto the input stream and rescanned.

M4 makes available the following built-in macros. They may be redefined, but once this is done the original meaning is lost. Their values are null unless otherwise stated.

define — The second argument is installed as the value of the macro whose name is the first argument. Each occurrence of **$***n* in the replacement text, where *n* is a digit, is replaced by the *n*-th argument. Argument **0** is the name of the macro; missing arguments are replaced by the null string; **$#** is replaced by the number of arguments; **$*** is replaced by a list of all the arguments separated by commas; **$@** is like **$***, but each argument is quoted (with the current quotation marks).

undefine — Removes the definition of the macro named in its argument.

defn — Returns the quoted definition of its argument(s). It is useful for renaming macros, especially built-ins.

pushdef — Like *define*, but saves any previous definition.

popdef — Removes current definition of its argument(s), exposing the previous one if any.

ifdef — If the first argument is defined, the value is the second argument, otherwise the third. If there is no third argument, the value is null. The word XENIX is predefined in *M4*.

shift — Returns all but its first argument. The other arguments are quoted and pushed back with commas in between. The quoting nullifies the effect of the extra scan that will subsequently be performed.

changequote — Changes quotation marks to the first and second arguments. The symbols may be up to five characters long. *Changequote* without arguments restores the original values (i.e., ` ').

changecom — Changes left and right comment markers from the default # and newline. With no arguments, the comment mechanism is effectively disabled. With one argument, the left marker becomes the argument and the right marker becomes newline. With two arguments, both markers are affected. Comment markers may be up to five characters long.

divert — *M4* maintains 10 output streams, numbered 0-9. The final output is the concatenation of the streams in numerical order; initially stream 0 is the current stream. The *divert* macro changes the current output stream to its (digit-string) argument. Output diverted to a stream other than 0 through 9 is discarded.

undivert — Causes immediate output of text from diversions named as arguments, or all diversions if no argument. Text may be undiverted into another diversion. Undiverting discards the diverted text.

divnum — Returns the value of the current output stream.

dnl — Reads and discards characters up to and including the next newline.

ifelse — Has three or more arguments. If the first argument is the same string as the second, then the value is the third argument. If not, and if there are more than four arguments, the process is repeated with arguments 4, 5, 6 and 7. Otherwise, the value is either the fourth string, or if it is not present, null.

incr — Returns the value of its argument incremented by 1. The value of the argument is calculated by interpreting an initial digit-string as a decimal number.

decr — Returns the value of its argument decremented by 1.

eval — Evaluates its argument as an arithmetic expression, using 32-bit arithmetic. Operators include +, - , *, /, % ^ (exponentiation), bitwise &, |, ^, and ~; relationals; parentheses. Octal and hex numbers may be specified as in C. The second argument specifies the

radix for the result; the default is **10**. The third argument may be used to specify the minimum number of digits in the result.

len	Returns the number of characters in its argument.
index	Returns the position in its first argument where the second argument begins (zero origin), or **– 1** if the second argument does not occur.
substr	Returns a substring of its first argument. The second argument is a zero origin number selecting the first character; the third argument indicates the length of the substring. A missing third argument is taken to be large enough to extend to the end of the first string.
translit	Transliterates the characters in its first argument from the set given by the second argument to the set given by the third. No abbreviations are permitted.
include	Returns the contents of the file named in the argument.
sinclude	Identical to *include*, except that it says nothing if the file is inaccessible.
syscmd	Executes the XENIX command given in the first argument. No value is returned.
sysval	Is the return code from the last call to *syscmd*.
maketemp	Fills in a string of XXXXX in its argument with the current process ID.
m4exit	Causes immediate exit from *m4*. Argument **1**, if given, is the exit code; the default is **0**.
m4wrap	Argument 1 will be pushed back at final EOF; example: m4wrap(`cleanup( ) ')
errprint	Prints its argument on the diagnostic output file.
dumpdef	Prints current names and definitions, for the named items, or for all if no arguments are given.
traceon	With no arguments, turns on tracing for all macros (including built-ins). Otherwise, turns on tracing for named macros.
traceoff	Turns off trace globally and for any macros specified. Macros specifically traced by *traceon* can be untraced only by specific calls to *traceoff*.

Name

make – Maintains, updates, and regenerates groups of programs.

Syntax

make [– **f** makefile] [– **p**] [– **i**] [– **k**] [– **s**] [– **r**] [– **n**] [– **b**] [– **e**] [– **t**] [– **q**] [– **d**] [names]

Description

The following is a brief description of all options and some special names:

– **f** *makefile* Description filename. *Makefile* is assumed to be the name of a description file. A filename of – denotes the standard input. The contents of *makefile* override the built-in rules if they are present.

– **p** Prints out the complete set of macro definitions and target descriptions.

– **i** Ignores error codes returned by invoked commands. This mode is entered if the fake target name **.IGNORE** appears in the description file.

– **k** Abandons work on the current entry, but continues on other branches that do not depend on that entry.

– **s** Silent mode. Does not print command lines before executing. This mode is also entered if the fake target name **.SILENT** appears in the description file.

– **r** Does not use the built-in rules.

– **n** No execute mode. Prints commands, but does not execute them. Even lines beginning with an @ are printed.

– **b** Compatibility mode for old makefiles.

– **e** Environment variables override assignments within makefiles.

– **t** Touches the target files (causing them to be up-to-date) rather than issues the usual commands.

– **d** Debug mode. Prints out detailed information on files and times examined.

– q	Question. The *make* command returns a zero or nonzero status code depending on whether the target file is or is not up-to-date.
.DEFAULT	If a file must be made but there are no explicit commands or relevant built-in rules, the commands associated with the name **.DEFAULT** are used if it exists.
.PRECIOUS	Dependents of this target will not be removed when quit or interrupt are hit.
.SILENT	Same effect as the – s option.
.IGNORE	Same effect as the – i option.

Make executes commands in *makefile* to update one or more target *names*. *Name* is typically a program. If no – f option is present, **makefile**, **Makefile**, **s.makefile**, and **s.Makefile** are tried in order. If *makefile* is – , the standard input is taken. More than one – f makefile argument pair may appear.

Make updates a target only if it depends on files that are newer than the target. All prerequisite files of a target are added recursively to the list of targets. Missing files are deemed to be out of date.

Makefile contains a sequence of entries that specify dependencies. The first line of an entry is a blank-separated, nonnull list of targets, then a :, then a (possibly null) list of prerequisite files or dependencies. Text following a ; and all following lines that begin with a tab are shell commands to be executed to update the target. The first line that does not begin with a tab or # begins a new dependency or macro definition. Shell commands may be continued across lines with the <backslash><newline> sequence. (#) and newline surround comments.

The following *makefile* says that **pgm** depends on two files **a.o** and **b.o**, and that they in turn depend on their corresponding source files (**a.c** and **b.c**) and a common file **incl.h**:

```
pgm: a.o b.o
     cc a.o b.o – o pgm
a.o: incl.h a.c
     cc – c a.c
b.o: incl.h b.c
     cc – c b.c
```

Command lines are executed one at a time, each by its own shell. A line is printed when it is executed unless the – s option is present, or the entry .SILENT: is in *makefile*, or unless the first character of the command is @ . The – n option specifies printing without execution; however, if the command line has the string **$(MAKE)** in it, the

line is always executed (see discussion of the **MAKEFLAGS** macro under *Environment*). The – **t** (touch) option updates the modified date of a file without executing any commands.

Commands returning nonzero status normally terminate *make*. If the – **i** option is present, or the entry **.IGNORE:** appears in *makefile*, or if the line specifying the command begins with <tab><hyphen>, the error is ignored. If the – **k** option is present, work is abandoned on the current entry, but continues on other branches that do not depend on that entry.

The – **b** option allows old makefiles (those written for the old version of *make*) to run without errors. The difference between the old version of *make* and this version is that this version requires all dependency lines to have a (possibly null) command associated with them. The previous version of *make* assumed if no command was specified explicitly that the command was null.

Interrupt and quit cause the target to be deleted unless the target depends on the special name **.PRECIOUS**.

Environment

The environment is read by *make*. All variables are assumed to be macro definitions and processed as such. The environment variables are processed before any makefile and after the internal rules; thus, macro assignments in a makefile override environment variables. The – **e** option causes the environment to override the macro assignments in a makefile.

The **MAKEFLAGS** environment variable is processed by *make* as containing any legal input option (except – **f**, – **p**, and – **d**) defined for the command line. Further, upon invocation, *make* "invents" the variable if it is not in the environment, puts the current options into it, and passes it on to invocations of commands. Thus, **MAKEFLAGS** always contains the current input options. This proves very useful for "super-makes". In fact, as noted above, when the – **n** option is used, the command **$(MAKE)** is executed anyway; hence, one can perform a **make** – **n** recursively on a whole software system to see what would have been executed. This is because the – **n** is put in **MAKEFLAGS** and passed to further invocations of **$(MAKE)**. This is one way of debugging all of the makefiles for a software project without actually doing anything.

Macros

Entries of the form *string1* = *string2* are macro definitions. Subsequent appearances of $(*string1*[:*subst1*=[*subst2*]]) are replaced by *string2*. The parentheses are optional if a single character macro name is used and there is no substitute sequence. The optional :*subst1*=*subst2* is a substitute sequence. If it is specified, all nonoverlapping occurrences of *subst1* in the named macro are replaced by

subst2. Strings (for the purposes of this type of substitution) are delimited by blanks, tabs, newline characters, and beginnings of lines. An example of the use of the substitute sequence is shown under *Libraries*.

Internal Macros

There are five internally maintained macros which are useful for writing rules for building targets:

$*** The macro **$* stands for the filename part of the current dependent with the suffix deleted. It is evaluated only for inference rules.

$@** The **$@ macro stands for the full target name of the current target. It is evaluated only for explicitly named dependencies.

$<** The **$< macro is only evaluated for inference rules or the **.DEFAULT** rule. It is the module which is out of date with respect to the target (i.e., the "manufactured" dependent filename). Thus, in the .c.o rule, the **$<** macro would evaluate to the .c file. An example for making optimized .o files from .c files is:

```
.c.o:
    cc - c - O $*.c
```

or:

```
.c.o:
    cc - c - O $<
```

$?** The **$? macro is evaluated when explicit rules from the makefile are evaluated. It is the list of prerequisites that are out of date with respect to the target; essentially, those modules which must be rebuilt.

$% The **$%** macro is only evaluated when the target is an archive library member of the form **lib(file.o)**. In this case, **$@** evaluates to **lib** and **$%** evaluates to the library member, **file.o**.

Four of the five macros can have alternative forms. When an upper case **D** or **F** is appended to any of the four macros the meaning is changed to "directory part" for **D** and "file part" for **F**. Thus, **$(@D)** refers to the directory part of the string **$@**. If there is no directory part ./ is generated. The only macro excluded from this alternative form is **$?**.

Suffixes

Certain names (for instance, those ending with .o) have default

dependents such as **.c**, **.s**, etc. If no update commands for such a file appear in *makefile*, and if a default dependent exists, that prerequisite is compiled to make the target. In this case, *make* has inference rules which allow building files from other files by examining the suffixes and determining an appropriate inference rule to use. The current default inference rules are:

```
.c .c~ .sh .sh~ .c.o .c~.o .c~.c .s.o .s~.o .y.o .y~.o .l.o .l~.o
.y.c .y~.c .l.c .c.a .c~.a .s~.a .h~.h
```

The internal rules for *make* are contained in the source file **rules.c** for the *make* program. These rules can be locally modified. To print out the rules compiled into the *make* on any machine in a form suitable for recompilation, the following command is used:

```
make - fp - 2>/dev/null </dev/null
```

The only peculiarity in this output is the **(null)** string which *printf*(S) prints when handed a null string.

A tilde in the above rules refers to an SCCS file (see *sccsfile*(F)). Thus, the rule **.c~.o** would transform an SCCS C source file into an object file (**.o**). Because the **s.** of the SCCS files is a prefix it is incompatible with *make*'s suffix point-of-view. Hence, the tilde is a way of changing any file reference into an SCCS file reference.

A rule with only one suffix (i.e. **.c:**) is the definition of how to build *x* from *x***.c**. In effect, the other suffix is null. This is useful for building targets from only one source file (e.g., shell procedures, simple C programs).

Additional suffixes are given as the dependency list for **.SUFFIXES**. Order is significant; the first possible name for which both a file and a rule exist is inferred as a prerequisite.

The default list is:

```
.SUFFIXES: .o .c .y .l .s
```

Here again, the above command for printing the internal rules will display the list of suffixes implemented on the current machine. Multiple suffix lists accumulate; **.SUFFIXES:** with no dependencies clears the list of suffixes.

Inference Rules

The first example can be done more briefly:

```
pgm: a.o b.o
    cc a.o b.o - o pgm
a.o b.o: incl.h
```

This is because *make* has a set of internal rules for building files. The user may add rules to this list by simply putting them in the *makefile*.

Certain macros are used by the default inference rules to permit the inclusion of optional matter in any resulting commands. For example, **CFLAGS**, **LFLAGS**, and **YFLAGS** are used for compiler options to *cc*(CP), *lex*(CP), and *yacc*(CP) respectively. Again, the previous method for examining the current rules is recommended.

The inference of prerequisites can be controlled. The rule to create a file with suffix .o from a file with suffix .c is specified as an entry with .c.o: as the target and no dependents. Shell commands associated with the target define the rule for making a .o file from a .c file. Any target that has no slashes in it and starts with a dot is identified as a rule and not as a true target.

Libraries

If a target or dependency name contains parentheses, it is assumed to be an archive library, the string within parentheses referring to a member within the library. Thus **lib(file.o)** and **$(LIB)(file.o)** both refer to an archive library which contains **file.o**. (This assumes the **LIB** macro has been previously defined.) The expression **$(LIB)(file1.o file2.o)** is not legal. Rules pertaining to archive libraries have the form *.XX*.**a** where the *XX* is the suffix from which the archive member is to be made. An unfortunate byproduct of the current implementation requires the *XX* to be different from the suffix of the archive member. Thus, one cannot have **lib(file.o)** depend upon **file.o** explicitly. The most common use of the archive interface follows. Here, we assume the source files are all C type source:

```
lib: lib(file1.o) lib(file2.o) lib(file3.o)
     @echo lib is now up to date
.c.a:
     $(CC) -c $(CFLAGS) $<
     ar rv $@ $*.o
     rm -f $*.o
```

In fact, the **.c.a** rule listed above is built into *make* and is unnecessary in this example. A more interesting, but more limited example of an archive library maintenance construction follows:

```
lib: lib(file1.o) lib(file2.o) lib(file3.o)
     $(CC) -c $(CFLAGS) $(?:.o=.c)
     ar rv lib $?
     rm $?   @echo lib is now up to date
.c.a:;
```

Here the substitution mode of the macro expansions is used. The **$?** list is defined to be the set of object filenames (inside **lib**) whose

C source files are out of date. The substitution mode translates the **.o** to **.c**. (Unfortunately, one cannot as yet transform to **.c~**) Note also, the disabling of the **.c.a:** rule, which would have created each object file, one by one. This particular construct speeds up archive library maintenance considerably. This type of construct becomes very cumbersome if the archive library contains a mix of assembly programs and C programs.

Files

[Mm]akefile

s.[Mm]akefile

See Also

sh(C)

Notes

Some commands return nonzero status inappropriately; use **– i** to overcome the difficulty. Commands that are directly executed by the shell, notably *cd*(C), are ineffectual across newlines in *make*. The syntax **(lib(file1.o file2.o file3.o)** is illegal. You cannot build **lib(file.o)** from file.o. The macro **$(a:.o=.c~)** is not available.

Name

mkstr – Creates an error message file from C source.

Syntax

mkstr [–] messagefile prefix file ...

Description

Mkstr is used to create files of error messages. Its use can make programs with large numbers of error diagnostics much smaller, and reduce system overhead in running the program as the error messages do not have to be constantly swapped in and out.

Mkstr will process each specified *file*, placing a massaged version of the input file in a file whose name consists of the specified *prefix* and the original name. The optional dash (–) causes the error messages to be placed at the end of the specified message file for recompiling part of a large *mkstr*ed program.

A typical *mkstr* command line is

```
mkstr pistrings xx *.c
```

This command causes all the error messages from the C source files in the current directory to be placed in the file *pistrings* and processed copies of the source for these files to be placed in files whose names are prefixed with *xx*.

To process the error messages in the source to the message file, *mkstr* keys on the string 'error(" in the input stream. Each time it occurs, the C string starting at the '"' is placed in the message file followed by a null character and a newline character; the null character terminates the message so it can be easily used when retrieved, the newline character makes it possible to sensibly *cat* the error message file to see its contents. The massaged copy of the input file then contains a *lseek* pointer into the file which can be used to retrieve the message. For example, the command changes

```
error("Error on reading", a2, a3, a4);
```

into

```
error(m, a2, a3, a4);
```

where *m* is the seek position of the string in the resulting error message file. The programmer must create a routine *error* which opens the message file, reads the string, and prints it out. The following example illustrates such a routine.

Example

```
char    efilname[] = "/usr/lib/pi_strings";
int     efil = -1;

error(a1, a2, a3, a4)
{
        char buf[256];

        if (efil < 0) {
                efil = open(efilname, 0);
                if (efil < 0) {
                        perror(efilname);
                        exit(C);
                }
        }
        if (lseek(efil, (long) a1, 0) ||read(efil, buf, 256) <= 0)
                goto oops;
        printf(buf, a2, a3, a4);
}
```

See Also

lseek(S), xstr(CP)

Credit

This utility was developed at the University of California at Berkeley and is used with permission.

Notes

All the arguments except the name of the file to be processed are unnecessary.

Name

nm – Prints name list.

Syntax

nm [– **acgnoOprsSuv**] [+ **offset**] [file ...]

Description

Nm prints the name list (symbol table) of each object *file* in the argument list. If an argument is an archive, a listing for each object file in the archive will be produced. If no *file* is given, the symbols in **a.out** are listed.

Each symbol name is preceded by its value in hexadecimal (blanks if undefined) and one of the letters **U** (undefined), **A** (absolute), **T** (text segment symbol), **D** (data segment symbol), **B** (bss segment symbol), S (segment name), **C** (common symbol), or **K** (8086 common segment). If the symbol table is in segmented format, symbol values are displayed as **segment:offset**. If the symbol is local (non-external) the type letter is in lowercase. The output is sorted alphabetically.

Options are:

– **a** Print only absolute symbols.

– **c** Print only C program symbols (symbols which begin with '_') as they appeared in the C program.

– **g** Print only global (external) symbols.

– **n** Sort numerically rather than alphabetically.

– **o** Prepend file or archive element name to each output line rather than only once.

– **O**
Print symbol values in octal.

– **p** Don't sort; print in symbol-table order.

– **r** Sort in reverse order.

– **s** Sort by size of symbol and display each symbol's size instead of value. The last symbol in each text or data segment ***may*** be assigned a size of **0**. This option implies the – **i** and – **n** options.

– S Switch the display format. If the symbol table is in segmented format, print values in non-segmented format. If not segmented, print values in segmented format.

– u Print only undefined symbols.

– v Also describe the object file and symbol table format.

Files

a.out Default input file

See Also

ar(CP), ar(F), a.out(F)

Name

prof – Displays profile data.

Syntax

prof [– a] [– l] [file]

Description

Prof interprets the file *mon.out* produced by the *monitor* subroutine. Under default modes, the symbol table in the named object file (*a.out* default) is read and correlated with the *mon.out* profile file. For each external symbol, the percentage of time spent executing between that symbol and the next is printed (in decreasing order), together with the number of times that routine was called and the number of milliseconds per call.

If the – a option is used, all symbols are reported rather than just external symbols. If the – l option is used, the output is listed by symbol value rather than decreasing percentage.

To cause calls to a routine to be tallied, the – p option of *cc* must have been given when the file containing the routine was compiled. This option also arranges for the *mon.out* file to be produced automatically.

Files

mon.out For profile

a.out For namelist

See Also

monitor(S), profil(S), cc(CP)

Notes

Beware of quantization errors.

If you use an explicit call to *monitor*(S) you will need to make sure that the buffer size is equal to or smaller than the program size.

Name

prs – Prints an SCCS file.

Syntax

prs [– d[dataspec]] [– r[SID]] [– e] [– l] [– a] files

Description

Prs prints, on the standard output, all or part of an SCCS file (see *sccsfile*(F)) in a user supplied format. If a directory is named, *prs* behaves as though each file in the directory were specified as a named file, except that nonSCCS files (last component of the pathname does not begin with s.), and unreadable files are silently ignored. If a name of – is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file or directory to be processed; nonSCCS files and unreadable files are silently ignored.

Arguments to *prs*, which may appear in any order, consist of options, and filenames.

All the described options apply independently to each named file:

– d[*dataspec*] Used to specify the output data specification. The *dataspec* is a string consisting of SCCS file *data keywords* (see *Data Keywords*) interspersed with optional user-supplied text.

– r[*SID*] Used to specify the *SCCS IDentification* (SID) string of a delta for which information is desired. If no SID is specified, the SID of the most recently created delta is assumed.

– e Requests information for all deltas created *earlier* than and including the delta designated via the – r option.

– l Requests information for all deltas created *later* than and including the delta designated via the – r option.

– a Requests printing of information for both removed, i.e., delta type = *R*, (see *rmdel*(CP)) and existing, i.e., delta type = *D*, deltas. If the – a option is not specified, information for existing deltas only is provided.

Data Keywords

Data keywords specify which parts of an SCCS file are to be retrieved and output. All parts of an SCCS file (see *sccsfile*(F)) have an associated data keyword. There is no limit on the number of times a data keyword may appear in a *dataspec*.

The information printed by *prs* consists of the user-supplied text and appropriate values (extracted from the SCCS file) substituted for the recognized data keywords in the order of appearance in the *dataspec*. The format of a data keyword value is either simple, in which keyword substitution is direct, or multiline, in which keyword substitution is followed by a carriage return.

User-supplied text is any text other than recognized data keywords. A tab is specified by \t and carriage return/newline is specified by \n.

TABLE 1. SCCS Files Data Keywords

Keyword	*Data Item*	*File Section*	*Value*	*Format*
:Dt:	Delta information	Delta Table	See below•	S
:DL:	Delta line statistics	"	:Li:/:Ld:/:Lu:	S
:Li:	Lines inserted by Delta	"	nnnnn	S
:Ld:	Lines deleted by Delta	"	nnnnn	S
:Lu:	Lines unchanged by Delta	"	nnnnn	S
:DT:	Delta type	"	*D* or *R*	S
:I:	SCCS ID string (SID)	"	:R:.:L:.:B:.:S:	S
:R:	Release number	"	nnnn	S
:L:	Level number	"	nnnn	S
:B:	Branch number	"	nnnn	S
:S:	Sequence number	"	nnnn	S
:D:	Date Delta created	"	:Dy:/:Dm:/:Dd:	S
:Dy:	Year Delta created	"	nn	S
:Dm:	Month Delta created	"	nn	S
:Dd:	Day Delta created	"	nn	S
:T:	Time Delta created	"	:Th::Tm::Ts:	S
:Th:	Hour Delta created	"	nn	S
:Tm:	Minutes Delta created	"	nn	S
:Ts:	Seconds Delta created	"	nn	S
:P:	Programmer who created Delta	"	logname	S
:DS:	Delta sequence number	"	nnnn	S
:DP:	Predecessor Delta seq-no.	"	nnnn	S
:DI:	Seq-no. of deltas incl., excl., ignored	"	:Dn:/:Dx:/:Dg:	S
:Dn:	Deltas included (seq #)	"	:DS: :DS: ...	S
:Dx:	Deltas excluded (seq #)	"	:DS: :DS: ...	S
:Dg:	Deltas ignored (seq #)	"	:DS: :DS: ...	S
:MR:	MR numbers for delta	"	text	M
:C:	Comments for delta	"	text	M
:UN:	User names	User Names	text	M
:FL:	Flag list	Flags	text	M
:Y:	Module type flag	"	text	S
:MF:	MR validation flag	"	*yes* or *no*	S
:MP:	MR validation pgm name	"	text	S
:KF:	Keyword error/warning flag	"	*yes* or *no*	S
:BF:	Branch flag	"	*yes* or *no*	S
:J:	Joint edit flag	"	*yes* or *no*	S
:LK:	Locked releases	"	:R: ...	S
:Q:	User defined keyword	"	text	S
:M:	Module name	"	text	S
:FB:	Floor boundary	"	:R:	S
:CB:	Ceiling boundary	"	:R:	S
:Ds:	Default SID	"	:I:	S
:ND:	Null delta flag	"	*yes* or *no*	S
:FD:	File descriptive text	Comments	text	M
:BD:	Body	Body	text	M
:GB:	Gotten body	"	text	M
:W:	A form of *what*(C) string	N/A	:Z::M:\t:I:	S
:A:	A form of *what*(C) string	N/A	:Z::Y: :M: :I::Z:	S
:Z:	*what*(C) string delimiter	N/A	@(#)	S
:F:	SCCS filename	N/A	text	S
:PN:	SCCS file pathname	N/A	text	S

• :Dt: = :DT: :I: :D: :T: :P: :DS: :DP:

Examples

The following:

```
prs – d"Users and/or user IDs for :F: are:\n:UN:" s.file
```

may produce on the standard output:

```
Users and/or user IDs for s.file are:
xyz
131
abc
```

```
prs – d"Newest delta for pgm :M:: :I: Created :D: By :P:" – r
s.file
```

may produce on the standard output:

```
Newest delta for pgm main.c: 3.7 Created 77/12/1 By cas
```

As a *special case:*

```
prs s.file
```

may produce on the standard output:

```
D 1.1 77/12/1 00:00:00 cas 1 000000/00000/00000
MRs:
bl78-12345
bl79-54321
COMMENTS:
this is the comment line for s.file initial delta
```

for each delta table entry of the "D" type. The only option allowed to be used with the *special case* is the – a option.

Files

/tmp/pr?????

See Also

admin(CP), delta(CP), get(CP), help(CP), sccsfile(F)

Diagnostics

Use *help*(CP) for explanations.

Name

ranlib – Converts archives to random libraries.

Syntax

ranlib arch1 arch2 ...

Description

Ranlib converts each archive to a form which can be utilized more rapidly by the linker, by adding a table of contents named __.SYMDEF to the beginning of the archive.

See Also

ld(CP), ar(CP), copy(C), settime(C)

Notes

The user must make sure that the most recent library versions have been processed with *ranlib* before linking. If this is not done, *ld*(CP) cannot create executable programs using these libraries. Sufficient temporary file space must be available in /tmp.

Name

ratfor – Converts Rational FORTRAN into standard FORTRAN.

Syntax

ratfor [option ...] [filename ...]

Description

Ratfor converts a rational dialect of FORTRAN into ordinary irrational FORTRAN. *Ratfor* provides control flow constructs essentially identical to those in C:

```
statement grouping:
    { statement; statement; statement }

decision-making:
    if (condition) statement [ else statement ]
    switch (integer value) {
            case integer:    statement
            ...
            [ default: ]     statement
    }

loops:
    while (condition) statement
    for (expression; condition; expression) statement
    do limits statement
    repeat statement [ until (condition) ]
    break [n]
    next [n]
```

and some additional syntax to make programs easier to read and write:

```
Free form input:
    multiple statements/line; automatic continuation

Comments:
    # this is a comment

Translation of relationals:
    >, >=, etc., become .GT., .GE., etc.

Return (expression)
    returns expression to caller from function

Define:
    define name replacement
```

Include:
include filename

The option – h causes quoted strings to be turned into 27H constructs. – C copies comments to the output, and attempts to format it neatly. Normally, continuation lines are marked with an & in column 1; the option – 6x makes the continuation character x and places it in column 6.

Notes

This program translates its input into C source code, which in segmented programming environments, is suitable for compiling as a small model program only (see *cc*(CP)).

Name

regcmp – Compiles regular expressions.

Syntax

regcmp [–] file

Description

Regcmp, in most cases, precludes the need for calling *regcmp* (see *regex*(S)) from C programs. This saves on both execution time and program size. The command *regcmp* compiles the regular expressions in *file* and places the output in *file*.**i**. If the – option is used, the output will be placed in *file*.**c**. The format of entries in *file* is a name (C variable) followed by one or more blanks followed by a regular expression enclosed in double quotation marks. The output of *regcmp* is C source code. Compiled regular expressions are represented as **extern char** vectors. *File*.**i** files may thus be *included* into C programs, or *file*.**c** files may be compiled and later loaded. In the C program which uses the *regcmp* output, *regex*(*abc*,*line*) applies the regular expression named *abc* to *line*. Diagnostics are self-explanatory.

Examples

```
name    "([A- Za- z][A- Za- z0- 9_]*)$0"

telno   "\({0,1}([2- 9][01][1- 9])$0\){0,1} *"
        "([2- 9][0- 9]{2})$1[ - ]{0,1}"
        "([0- 9]{4})$2"
```

In the C program that uses the *regcmp* output,

```
regex(telno, line, area, exch, rest)
```

will apply the regular expression named *telno* to *line*.

See Also

regex(S)

Notes

This program translates its input into C source code, which in segmented programming environments, is suitable for compiling as a small model program only (see *cc*(CP)).

Name

rmdel – Removes a delta from an SCCS file.

Syntax

rmdel – r*SID files*

Description

Rmdel removes the delta specified by the *SID* from each named SCCS file. The delta to be removed must be the newest (most recent) delta in its branch in the delta chain of each named SCCS file. In addition, the SID specified must *not* be that of a version being edited for the purpose of making a delta. That is, if a *p-file* exists for the named SCCS file, the SID specified must *not* appear in any entry of the *p-file*(see *get*(CP)).

If a directory is named, *rmdel* behaves as though each file in the directory were specified as a named file, except that nonSCCS files (last component of the pathname does not begin with **s.**) and unreadable files are silently ignored. If a name of – is given, the standard input is read; each line of the standard input is taken to be the name of an SCCS file to be processed; nonSCCS files and unreadable files are silently ignored.

Files

x-file See *delta*(CP)

z-file See *delta*(CP)

See Also

delta(CP), get(CP), help(CP), prs(CP), sccsfile(F)

Diagnostics

Use *help*(CP) for explanations.

Name

sact – Prints current SCCS file editing activity.

Syntax

sact files

Description

Sact informs the user of any impending deltas to a named SCCS file. This situation occurs when *get*(CP) with the – e option has been previously executed without a subsequent execution of *delta*(CP). If a directory is named on the command line, ***sact*** behaves as though each file in the directory were specified as a named file, except that nonSCCS files and unreadable files are silently ignored. If a name of – is given, the standard input is read with each line being taken as the name of an SCCS file to be processed.

The output for each named file consists of five fields separated by spaces.

Field 1	Specifies the SID of a delta that currently exists in the SCCS file to which changes will be made to make the new delta
Field 2	Specifies the SID for the new delta to be created
Field 3	Contains the logname of the user who will make the delta i.e., executed a *get* for editing
Field 4	Contains the date that **get** – **e** was executed
Field 5	Contains the time that **get** – **e** was executed

See Also

delta(CP), get(CP), unget(CP)

Diagnostics

Use *help*(CP) for explanations.

Name

sccsdiff – Compares two versions of an SCCS file.

Syntax

sccsdiff – rSID1 **–** rSID2 [**– p**] [**– sn**] files

Description

Sccsdiff compares two versions of an SCCS file and generates the differences between the two versions. Any number of SCCS files may be specified, but arguments apply to all files.

– r*SID?* *SID1* and *SID2* specify the deltas of an SCCS file that are to be compared. Versions are passed to *bdiff*(C) in the order given.

– **p** Pipe output for each file through *pr*(C).

– **s***n* *n* is the file segment size that *bdiff* will pass to *diff*(C). This is useful when *diff* fails due to a high system load.

Files

/tmp/get????? Temporary files

See Also

bdiff(C), get(CP), help(CP), pr(C)

Diagnostics

file: *No differences* If the two versions are the same.

Use *help*(CP) for explanations.

Name

size – Prints the size of an object file.

Syntax

size [object ...]

Description

Size prints the (decimal) number of bytes required by the text, data, and bss portions, and their sum in decimal and hexadecimal, of each object-file argument. If no file is specified, **a.out** is used.

See Also

a.out(F)

Name

spline – Interpolates smooth curve.

Syntax

spline [option] ...

Description

Spline takes pairs of numbers from the standard input as abcissas and ordinates of a function. It produces a similar set, which is approximately equally spaced and includes the input set, on the standard output. The cubic spline output has two continuous derivatives, and enough points to look smooth when plotted.

The following options are recognized, each as a separate argument.

– **a** Supplies abscissas automatically (they are missing from the input); spacing is given by the next argument, or is assumed to be 1 if next argument is not a number.

– **k** The constant k used in the boundary value computation

$$y_0'' = k y_1' \ , \ \ldots \ , \ y_n'' = k y_{n-1}'$$

is set by the next argument. By default $k = 0$.

– **n** Spaces output points so that approximately n intervals occur between the lower and upper x limits. (Default $n = 100$.)

– **p** Makes output periodic, i.e. matches derivatives at ends. First and last input values should normally agree.

– **x** Next 1 (or 2) arguments are lower (and upper) x limits. Normally these limits are calculated from the data. Automatic abcissas start at lower limit (default 0).

Diagnostics

When data is not strictly monotone in x, *spline* reproduces the input without interpolating extra points.

Notes

A limit of 1000 input points is silently enforced.

Name

stackuse – Determines stack requirements for C programs.

Syntax

stackuse [– m *startsym*] [– r *fakeref*] [– s *libstack*] [– l *library*] [file ...]

Description

Stackuse determines the stack requirements of one or more C language programs. It displays the name of the *main* routine in a file, its stack requirements in bytes, and the number of recursive routines. All command line switches are optional.

– m*startsym* Print only the specified start ("main") symbol. If this option is not specified all start symbols (those which are not called by anybody) will be printed.

– r*fakeref* Uses the named file *fakeref* as a fake references file. The format is: *parent child* . The special *parent* .LEAF is a meta-parent meaning all leaf nodes.

– s*libstack* Uses the named file as library of costs for external routines. The format is: *subr stack* . The special *subr* .UNDEF is a meta-subroutine meaning all undefined routines.

– l*library* Uses a system-provided libstack for a standard libraries. e.g. – lc, – ll, – ly.

– a Print data for all symbols, not just start symbols.

The – r and – s options may be repeated an arbitrary number of times. The effect is additive rather than destructive. In the case of duplicate definitions, the first is used.

Lines of the – r and – s files which begin with a pound sign (#) are treated as comments, and otherwise are ignored.

Diagnostics

Usage (fatal).

Redefinitions in – r, – s files, or in the source (warning).

Presence of routines for which no stack value is provided (warning).

Files

/usr/lib/stackuse/*	Passes, libraries
/tmp/*	Temporaries used by passes.

Notes

For the **libstack** and **fakeref** files, a comment character (#) is used.

STRINGS (CP)

Name

strings – Finds the printable strings in an object file.

Syntax

strings [–] [– o] [– *number*] file ...

Description

Strings looks for ASCII strings in a binary file. A string is any sequence of four or more printing characters ending with a newline or a null character. Unless the – flag is given, *strings* only looks in the initialized data space of object files. If the – o flag is given, then each string is preceded by its decimal offset in the file. If the – *number* flag is given then *number* is used as the minimum string length rather than 4.

Strings is useful for identifying random object files and many other things.

See Also

hd(C), od(C)

Credit

This utility was developed at the University of California at Berkeley and is used with permission.

Name

strip – Removes symbols and relocation bits.

Syntax

strip [dehrsStx] file ...

Description

Strip removes selcted parts of an object file, including the header, text, data, relocation records, and symbol table. Strip works directly on the named files; nothing is written to the standard output.

Strip is typically used to remove symbol table and relocation information from a file after debugging has been completed. It also is useful for creating a compact namelist file in which text and data have been removed.

– **d** Strip data and the data relocation records.

– **e** Strip the extended header.

– **h** Strip the header and extended header.

– **r** Strip all relocation records except the **x.out** short form.

– **s** Strip the symbol table.

– **S** Strip the segment table.

– **t** Strip text and the text relocation records.

– **x** Strip all relocation records.

Strip has the same effect as the – **s** option of *ld*. If no options are given, the – **r** and – **s** options are implied.

Although ***strip*** can be used to remove an **x.out** header from an 8086 relocatable file, it cannot be used to remove run-time relocation records.

Files

/tmp/s* Temporary file

See Also

ld(CP), a.out(F)

Name

time – Times a command.

Syntax

time command

Description

The given *command* is executed; after it is complete, *time* prints the elapsed time during the command, the time spent in the system, and the time spent in execution of the command. Times are reported in seconds.

The times are printed on the standard error.

See Also

times(S)

Name

tsort – Sorts a file topologically.

Syntax

tsort [file]

Description

Tsort produces on the standard output a totally ordered list of items consistent with a partial ordering of items mentioned in the input *file*. If no *file* is specified, the standard input is understood.

The input consists of pairs of items (nonempty strings) separated by blanks. Pairs of different items indicate ordering. Pairs of identical items indicate presence, but not ordering.

See Also

lorder(CP)

Diagnostics

Odd data: There is an odd number of fields in the input file.

Notes

The *sort* algorithm is quadratic, which can be slow if you have a large input list.

Name

unget – Undoes a previous get of an SCCS file.

Syntax

unget [– r*SID*] [– s] [– n] files

Description

Unget undoes the effect of a **get – e** done prior to creating the intended new delta. If a directory is named, ***unget*** behaves as though each file in the directory were specified as a named file, except that nonSCCS files and unreadable files are silently ignored. If a name of – is given, the standard input is read with each line being taken as the name of an SCCS file to be processed.

Options apply independently to each named file.

– r*SID* Uniquely identifies which delta is no longer intended. (This would have been specified by *get* as the "new delta".) The use of this option is necessary only if two or more versions of the same SCCS file have been retrieved for editing by the same person (login name). A diagnostic results if the specified *SID* is ambiguous, or if it is necessary and omitted on the command line.

– s Suppresses the printout, on the standard output, of the intended delta's *SID*.

– n Causes the retention of the file which would normally be removed from the current directory.

See Also

delta(CP), get(CP), sact(CP)

Diagnostics

Use *help*(CP) for explanations.

Name

uucp, uulog – Copies files from XENIX to XENIX.

Syntax

uucp [option] ... source-file ... destination-file

uulog [option] ...

Description

Uucp copies files named by the source-file arguments to the destination-file argument. A filename may be a pathname on your machine, or may have the form:

system-name!pathname

where "system-name" is taken from a list of system names which *uucp* knows about. Shell metacharacters ?*[] appearing in *pathname* will be expanded on the appropriate system.

Pathnames may be a a full pathname, or a pathname preceded by ~*user* where *user* is a user ID on the specified system and is replaced by that user's login directory. Anything else is prefixed by the current directory.

If the result is an erroneous pathname for the remote system the copy will fail. If the destination-file is a directory, the last part of the source-filename is used.

Uucp preserves execute permissions across the transmission and gives 0666 read and write permissions (see *chmod*(S)).

The following options are interpreted by *uucp*:

– d Makes all necessary directories for the file copy.

– c Uses the source file when copying out rather than copying the file to the spool directory.

– m Sends mail to the requester when the copy is complete.

Uulog maintains a summary log of *uucp* and *uux*(CP) transactions in the file /usr/spool/uucp/LOGFILE by gathering information from partial log files named /usr/spool/uucp/LOG.*.? . It removes the partial log files.

The options cause *uulog* to print logging information:

– s*sys*
Prints information about work involving system *sys*.

– u*user*
Prints information about work done for the specified *user*.

Files

/usr/spool/uucp Spool directory

/usr/lib/uucp/* Other data and program files

See Also

uux(CP), mail(C)

Warning

The domain of remotely accessible files can (and for obvious security reasons, usually should) be severely restricted. You will very likely not be able to fetch files by pathname; ask a responsible person on the remote system to send them to you. For the same reasons you will probably not be able to send files to arbitrary pathnames.

Notes

For security reasons, all files received by *uucp* should be owned by uucp.

The – m option will only work sending files or receiving a single file. Receiving multiple files specified by special shell characters ? *[] will not activate the – m option.

This version of *uucp* is based on a version 7 implementation of the program.

Name

uux – Executes command on remote XENIX.

Syntax

uux [–] command-string

Description

Uux will gather 0 or more files from various systems, execute a command on a specified system and send standard output to a file on a specified system.

The command-string is made up of one or more arguments that look like a shell command line, except that the command and filenames may be prefixed by system-name!. A null system-name is interpreted as the local system.

Filenames may be (1) a full pathname; (2) a pathname preceded by ~*xxx*; where *xxx* is a user ID on the specified system and is replaced by that user's login directory; or (3) anything else prefixed by the current directory.

The "– " option will cause the standard input to the *uux* command to be the standard input to the command-string.

For example, the command

uux "!diff usg!/usr/dan/f1 pwba!/a4/dan/f1 > !fi.diff"

will get the **f1** files from the usg and pwba machines, execute a *diff* command and put the results in **f1.diff** in the local directory.

Any special shell characters such as <>;| should be quoted either by quoting the entire command-string, or quoting the special characters as individual arguments.

Files

/usr/uucp/spool Spool directory

/usr/uucp/* Other data and programs

See Also

uucp(CP)

Warning

An installation may, and for security reasons generally will, limit the list of commands executable on behalf of an incoming request from *uux*. Typically, a restricted site will permit little other than the receipt of mail via *uux*.

Notes

Only the first command of a shell pipeline may have a system-name!. All other commands are executed on the system of the first command.

The shell metacharacter * will probably not perform as expected.

The shell tokens << and >> are not implemented.

There is no notification of denial of execution on the remote machine.

Name

val – Validates an SCCS file.

Syntax

val –

val [– s] [– r*SID*] [– m*name*] [– y*type*] files

Description

Val determines if the specified *file* is an SCCS file meeting the characteristics specified by the optional argument list. Arguments to *val* may appear in any order. The arguments consist of options, which begin with a – , and named files.

Val has a special argument, – , which causes reading of the standard input until an end-of-file condition is detected. Each line read is independently processed as if it were a command line argument list.

Val generates diagnostic messages on the standard output for each command line and file processed and also returns a single 8-bit code upon exit as described below.

The options are defined as follows. The effects of any option apply independently to each named file on the command line:

– s — The presence of this argument silences the diagnostic message normally generated on the standard output for any error that is detected while processing each named file on a given command line.

– r*SID* — The argument value *SID* (*S*CCS *ID*entification String) is an SCCS delta number. A check is made to determine if the *SID* is ambiguous (e. g., **r1** is ambiguous because it physically does not exist but implies **1.1**, **1.2**, etc. which may exist) or invalid (e. g., **r1.0** or **r1.1.0** are invalid because neither case can exist as a valid delta number). If the *SID* is valid and not ambiguous, a check is made to determine if it actually exists.

– m*name* — The argument value *name* is compared with the SCCS %M% keyword in *file*.

– y*type* — The argument value *type* is compared with the SCCS %Y% keyword in *file*.

The 8-bit code returned by *val* is a disjunction of the possible errors, i. e., can be interpreted as a bit string where (moving from left to right) set bits are interpreted as follows:

bit 0 = Missing file argument

bit 1 = Unknown or duplicate option

bit 2 = Corrupted SCCS file

bit 3 = Can't open file or file not SCCS

bit 4 = *SID* is invalid or ambiguous

bit 5 = *SID* does not exist

bit 6 = %Y%, – y mismatch

bit 7 = %M%, – m mismatch

Note that *val* can process two or more files on a given command line and in turn can process multiple command line (when reading the standard input). In these cases an aggregate code is returned; a logical OR of the codes generated for each command line and file processed.

See Also

admin(CP), delta(CP), get(CP), prs(CP)

Diagnostics

Use *help*(CP) for explanations.

Notes

Val can process up to 50 files on a single command line.

XREF (CP)

Name

xref – Cross-references C programs.

Syntax

xref [file ...]

Description

Xref reads the named *files* or the standard input if no file is specified and prints a cross reference consisting of lines of the form

identifier filename line numbers ...

Function definition is indicated by a plus sign (+) preceding the line number.

See Also

cref(CP)

Name

xstr – Extracts strings from C programs.

Syntax

xstr [– c] [–] [file]

Description

Xstr maintains a file *strings* into which strings in component parts of a large program are hashed. These strings are replaced with references to this common area. This serves to implement shared constant strings, most useful if they are also read-only.

The command

xstr – c name

will extract the strings from the C source in name, replacing string references by expressions of the form (&xstr[number]) for some number. An appropriate declaration of *xstr* is prepended to the file. The resulting C text is placed in the file *x.c*, to then be compiled. The strings from this file are placed in the *strings* data base if they are not there already. Repeated strings and strings which are suffices of existing strings do not cause changes to the data base.

After all components of a large program have been compiled, a file *xs.c* declaring the common *xstr* space can be created by a command of the form

xstr -c name1 name2 name3 ...

This *xs.c* file should then be compiled and loaded with the rest of the program. If possible, the array can be made read-only (shared) saving space and swap overhead.

Xstr can also be used on a single file. A command

xstr name

creates files *x.c* and *xs.c* as before, without using or affecting any *strings* file in the same directory.

It may be useful to run *xstr* after the C preprocessor if any macro definitions yield strings or if there is conditional code which contains strings which may not, in fact, be needed. *Xstr* reads from its standard input when the argument – is given. An appropriate command sequence for running *xstr* after the C preprocessor is:

```
cc – E name.c |xstr – c –
cc – c x.c
mv x.o name.o
```

Xstr does not touch the file *strings* unless new items are added, thus *make* can avoid remaking *xs.o* unless truly necessary.

Files

strings	Data base of strings
x.c	Massaged C source
xs.c	C source for definition of array "xstr"
/tmp/xs*	Temp file when "xstr name" doesn't touch *strings*

See Also

mkstr(CP)

Credit

This utility was developed at the University of California at Berkeley and is used with permission.

Notes

If a string is a suffix of another string in the data base, but the shorter string is seen first by *xstr* , both strings will be placed in the data base when just placing the longer one there will do.

Name

yacc – Invokes a compiler-compiler.

Syntax

yacc [– vd] grammar

Description

Yacc converts a context-free grammar into a set of tables for a simple automaton which executes an LR(1) parsing algorithm. The grammar may be ambiguous; specified precedence rules are used to break ambiguities.

The output file, y.tab.c, must be compiled by the C compiler to produce a program *yyparse*. This program must be loaded with the lexical analyzer program, *yylex*, as well as *main* and *yyerror*, an error handling routine. These routines must be supplied by the user; *lex*(CP) is useful for creating lexical analyzers usable by *yacc*.

If the – v flag is given, the file y.output is prepared, which contains a description of the parsing tables and a report on conflicts generated by ambiguities in the grammar.

If the – d flag is used, the file y.tab.h is generated with the **#define** statements that associate the *yacc*-assigned "token codes" with the user-declared "token names". This allows source files other than y.tab.c to access the token codes.

Files

y.output	
y.tab.c	
y.tab.h	Defines for token names
yacc.tmp, yacc.acts	Temporary files
/usr/lib/yaccpar	Parser prototype for C programs

See Also

lex(CP)

Diagnostics

The number of reduce-reduce and shift-reduce conflicts is reported on the standard output; a more detailed report is found in the **y.output** file. Similarly, if some rules are not reachable from the start symbol, this is also reported.

Notes

Because filenames are fixed, at most one *yacc* process can be active in a given directory at a time.

This program translates its input into C source code, which in segmented programming environments, is suitable for compiling as a small model program only (see *cc*(CP)).

Index

Archives and libraries ..... **ar**
C language usage and syntax ..... **lint**
C program, formatting ..... **cb**
Commands, intersystem execution ..... **uux**
Compiler compiler ..... **yacc**
Debugger ..... **adb**
Error message file ..... **mkstr**
Execution, time ..... **time**
File, intersystem copy ..... **uucp**
Graphics, interpolating curves ..... **spline**
Lexical analyzers ..... **lex**
Macro processor **m4**
Object file, printable strings ..... **strings**
Object file, size ..... **size**
Object file, displaying ..... **hdr**
Object file, symbols and relocation ..... **strip**
Ordering relations ..... **lorder**
Program listing, cross-reference ..... **xref**
Program listing, cross-reference ..... **cref**
Program maintenance ..... **make**
Rational FORTRAN ..... **ratfor**
Regular expressions ..... **regcmp**
SCCS files, combining ..... **comb**
SCCS files, comments ..... **cdc**
SCCS files, comparing ..... **sccsdiff**
SCCS files, creating new versions ..... **delta**
SCCS files, editing ..... **sact**
SCCS files, printing ..... **prs**
SCCS files, removing ..... **rmdel**
SCCS files, restoring ..... **unget**
SCCS files, retrieving versions ..... **get**
SCCS files, creating and maintaining ..... **admin**
SCCS files, validating ..... **val**
SCCS, command help ..... **help**
Shell ..... **csh**
Sorting topologically ..... **tsort**
Stack requirements ..... **stackuse**
Standard input, reading strings ..... **gets**
Strings, extracting ..... **xstr**
System, XENIX configuration ..... **config**
Tags file ..... **ctags**
uulog command ..... **uucp**